arXiv:1603.09055v1 [cs.LO] 30 Mar 2016

Succinctness of Order-Invariant Logics on Depth-Bounded Structures*

Kord Eickmeyer[1], Michael Elberfeld[2], and Frederik Harwath[3]

[1]TU Darmstadt, Germany
[2]RWTH Aachen University, Germany
[3]Goethe University Frankfurt, Germany

March 31, 2016

Abstract

We study the expressive power and succinctness of order-invariant sentences of first-order (FO) and monadic second-order (MSO) logic on structures of bounded tree-depth. Order-invariance is undecidable in general and, thus, one strives for logics with a decidable syntax that have the same expressive power as order-invariant sentences. We show that on structures of bounded tree-depth, order-invariant FO has the same expressive power as FO. Our proof technique allows for a fine-grained analysis of the succinctness of this translation. We show that for every order-invariant FO sentence there exists an FO sentence whose size is elementary in the size of the original sentence, and whose number of quantifier alternations is linear in the tree-depth. We obtain similar results for MSO. It is known that the expressive power of MSO and FO coincide on structures of bounded tree-depth. We provide a translation from MSO to FO and we show that this translation is essentially optimal regarding the formula size. As a further result, we show that order-invariant MSO has the same expressive power as FO with modulo-counting quantifiers on bounded tree-depth structures.

1 Introduction

Understanding the *expressivity* of logics on finite structures—the question of which properties are definable in a certain logic—plays an important role in database and complexity theory. In the former, logics are used to formulate queries; in the latter, they describe computational problems. Moreover, besides just studying a logic's expressivity, understanding its *succinctness*—the question of how complex definitions of properties such as queries and problems must be—is a requirement towards (theoretical) expressivity results of (potential) practical importance. The present work studies the succinctness of first-order logic (FO) as well as its succinctness compared to extensions allowing for the use of a linear order and set quantifiers. This extends and refines recent studies on the expressivity of these logics [1, 8] on restricted classes of structures. The structures we consider have bounded tree-depth, which is a graph invariant that measures how far a graph is from being a star in a similar way as tree-width measures how far a graph is from being a tree. Our results are summarised by Figure 1.

*A preliminary version of this paper was presented at the MFCS 2014 conference [7].

$\varphi \in$	$<$-inv-FO	MSO	$<$-inv-MSO
$\psi \in$	FO	FO	FO+MOD
$\|\psi\|$	d-EXP(q)	d-EXP(q)	non-elementary
$\mathrm{qad}(\psi)$	$O(d)$	$O(d)$	$O(d)$

Figure 1: Summary of our results: A formula φ of quantifier rank q is translated into a formula ψ that is equivalent to φ on structures of tree-depth at most d.

In both database and complexity theory, one often assumes that structures come with a linear order and formulae are allowed to use this order as long as the properties defined by them do not depend on the concrete interpretation of the order in a structure. Such formulae are called *order-invariant*. Since testing order-invariance for given FO-formulae is undecidable in general, one tries to find logics that have the same expressive power as order-invariant formulae, but a decidable syntax. Several examples prove that order-invariant FO-formulae ($<$-inv-FO) are more expressive than FO-formulae without access to orders, cf. [18]. A common feature of these separating examples is that their Gaifman graphs contain large cliques, making them rather complicated from the point of view of graph structure theory.

For tree structures, on the other hand, [1] showed that the expressivity of FO and $<$-inv-FO coincide. Following this example, we show that on structures of tree-depth at most d each $<$-inv-FO-sentence can be translated to an FO-sentence whose size is d-fold exponential in the size of the original sentence (Theorem 4). The importance of the expressivity result is highlighted by the fact that order-invariance is undecidable even on structures of tree-depth at most 2 (Theorem 5).

A logic that is commonly studied from the perspectives of algorithm design and language theory is monadic second-order logic (MSO), which extends FO-formulae by the ability to quantify over sets of elements instead of just single elements. While it has a rich expressivity that exceeds that of FO already on word structures, the expressive powers of FO and MSO coincide on any class of structures whose tree-depth is bounded [8] by a constant d. We refine this by presenting a translation into FO-formulae of d-fold exponential size (Theorem 18). We prove that this translation is essentially optimal regarding the formula size (Theorem 19). Beside the succinctness results, we prove that $<$-inv-MSO has the same expressive power as FO+MOD, the extension of FO by arbitrary first-order modulo-counting quantifiers, for structures of bounded tree-depth (Theorem 14).

Our results also have implications for FO itself. They imply that the quantifier alternation hierarchy for FO of [3] collapses on structures of bounded tree-depth, whereas it is shown in [3] to be strict on trees of unbounded height. For structures of bounded tree-depth we are able to turn any FO-formula into a formula whose size is bounded by the quantifier depth of the original formula and whose quantifier alternation depth is bounded by a linear function in the tree-depth.

A recurring theme in the study of FO, MSO, and their variants is the question of which graph-theoretical properties can be defined using formulae of these logics. The main motivation behind these questions lies in the fact that access to certain tree-decompositions or embeddings of the structure can be used as a proof ingredient for translating formulae. Independent of the results stated above, we prove that, for structures of bounded tree-depth, it is possible to define tree-decompositions of bounded width *and height* in FO (cf. Section 6).

Proof techniques Our proofs use techniques from finite model theory, in particular interpretation arguments, logical types, and games. Compared to prior works like [8], we enrich the application of these techniques by a quantitative analysis, thereby obtaining succinct translations instead of just equal expressivity results. The proofs of [8] use an involved constructive variant of the Feferman–Vaught composition theorem, which complicates a straightforward analysis of the formula size in the translation from MSO to FO. We also use composition arguments, but we get along with an easier non-constructive variant. There is another proof of the result of [8] in [10], but it relies on involved combinatorial insights that seem unsuited for both a tight analysis of succinctness as well as an adaptation to the ordered setting.

The results of [1] about the expressivity of <-inv-FO on trees use automata-theoretic and algebraic methods. Since these methods seem unsuited to obtain succinct formula translations, we apply and develop techniques that are mainly based on using games: In order to translate <-inv-FO-sentences into FO-sentences, we first restrict our attention to a certain kind of linear ordering that is based on the logical types of recursively-defined substructures. Since the FO-type of ordered structures turns out to be FO-definable in the original (unordered) structures, we are able to prove a succinct translation from <-inv-FO to FO.

In order to translate <-inv-MSO-sentences into FO+MOD-sentences, the proof structure is similar, but we need to add a "pumping lemma" for <-inv-MSO, which proves the limited expressive power of <-inv-MSO on the recursively considered substructures.

Organisation of this paper The paper continues with a background section and, then, the results related to <-inv-FO, MSO, and <-inv-MSO are proved in Sections 3, 5, and 4, respectively. Tree-decompositions for structures of bounded tree-depth are handled in Section 6.

2 Background

In the present section, we review definitions and terms related to logical formulae and structures as well as the notion of tree-depth.

General notation The sets of natural numbers with and without 0 are denoted, respectively, by $\mathbb{N}$ and $\mathbb{N}^+$. Let $[i,j] := \{i,\dots,j\}$ for all $i,j \in \mathbb{N}$ with $i \leq j$, and let $[j] := [1,j]$. We define the *d-fold exponential function* $d\text{-exp}(n)$ recursively by $0\text{-exp}(n) := n$, and $(d+1)\text{-exp}(n) := 2^{d\text{-exp}(n)}$. The *class of functions that grow at most d-fold exponentially* is $d\text{-EXP} := \{f : \mathbb{N} \to \mathbb{N} \mid f(n) \leq d\text{-exp}(n^c)$ for some $c \in \mathbb{N}$ and all $n > c\}$. If we say that a relation is an *order*, we implicitly assume that it is linear. Thus an order is an antisymmetric, transitive, reflexive and total binary relation.

Logic For a reference on notation and standard methods in finite model theory, we refer to the book of [13]. We denote structures by Fraktur letters $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}, \dots$ and their universes by the corresponding latin letters $A, B, C, \dots$. Besides the standard logics FO and MSO, we also consider the logic FO+MOD that is obtained from FO by allowing the use of *modulo-counting quantifiers* $\exists^{i \,(\mathrm{mod}\ p)}$ for each $i \in \mathbb{N}$ and $p \in \mathbb{N}^+$. The meaning of these quantifiers is that $\mathfrak{A} \models \exists^{i \,(\mathrm{mod}\ p)} x\, \varphi(x, \bar{y})$ iff $|\{b \in A : \mathfrak{A} \models \varphi(b, \bar{a})\}| \equiv i \,(\mathrm{mod}\ p)$, where $\mathfrak{A}$ is a structure and $\bar{a}$ is a tuple of its elements.

We write $\mathrm{qr}(\varphi)$ for the *quantifier rank* and $\|\varphi\|$ for the *size* (or *length*) of a formula φ. The *quantifier alternation depth* $\mathrm{qad}(\varphi)$ of a formula φ in *negation normal form* (NNF, i.e. all negations of φ occur directly in front of atomic formulae) is the maximum number of alternations between $\exists$-

and $\forall$-quantifiers on all directed paths in the syntax tree of φ. If φ is not in NNF, we first find an equivalent formula φ' in NNF using a fixed conversion procedure and then define $\mathrm{qad}(\varphi) := \mathrm{qad}(\varphi')$. If Φ is a set of formulae, we let $\|\Phi\| := \max_{\varphi\in\Phi} \|\varphi\|$ and $\mathrm{qad}(\Phi) := \max_{\varphi\in\Phi} \mathrm{qad}(\Phi)$.

For any logic $\mathrm{L} \in \{\mathrm{FO}, \mathrm{FO{+}MOD}, \mathrm{MSO}\}$ and $q \in \mathbb{N}$, we write $\mathfrak{A} \equiv^{\mathrm{L}}_q \mathfrak{B}$ for $q \in \mathbb{N}$ if σ-structures $\mathfrak{A}$ and $\mathfrak{B}$ satisfy the same $\mathrm{L}[\sigma]$-sentences of quantifier rank at most q. The $\equiv^{\mathrm{L}}_q$-equivalence class of $\mathfrak{A}$ is its (L, q)-*type* and is denoted by $\mathrm{tp}_{\mathrm{L},q}(\mathfrak{A})$. If the logic L has been fixed or the concrete logic is not important for the discussion, we omit it in this and similar notation.

For a signature σ, we denote by $\sigma^{\leq}$ the signature $\sigma \cup \{\leq\}$, where $\leq \notin \sigma$ is a binary relation symbol. An *ordered $\sigma^{\leq}$-structure* is a $\sigma^{\leq}$-structure $\mathfrak{A}$ where $\leq^{\mathfrak{A}}$ is an order on $\mathfrak{A}$. An *ordered expansion* $(\mathfrak{A}, \preceq)$ of a σ-structure $\mathfrak{A}$ is an expansion of $\mathfrak{A}$ to an ordered $\sigma^{\leq}$-structure. A sentence $\varphi \in \mathrm{FO}[\sigma^{\leq}]$ is *order-invariant* on a class $\mathcal{C}$ of structures if for all σ-structures $\mathfrak{A} \in \mathcal{C}$ and all ordered expansions $(\mathfrak{A}, \preceq_1)$ and $(\mathfrak{A}, \preceq_2)$ of $\mathfrak{A}$ we have $(\mathfrak{A}, \preceq_1) \models \varphi$ iff $(\mathfrak{A}, \preceq_2) \models \varphi$. If $\mathcal{C}$ is not otherwise stated, we assume $\mathcal{C}$ to be the class of all finite structures. The set of all order-invariant $\varphi \in \mathrm{FO}[\sigma, \leq]$ is denoted by $<$-inv-$\mathrm{FO}[\sigma]$, and for such a φ and a σ-structure $\mathfrak{A}$ we write $\mathfrak{A} \models_{\leq} \varphi$ if $(\mathfrak{A}, \preceq) \models \varphi$ for some (equivalently, for every) ordered expansion $(\mathfrak{A}, \preceq)$ of $\mathfrak{A}$; $<$-inv-MSO is defined in the same way.

The *restriction* of a binary relation R on a set M to a subset $N \subseteq M$ is the relation $R|_N := \{(x, y) \in R : x, y \in N\}$. Note that a substructure of an ordered structure is again an ordered structure. For two linear orders $\preceq_1$ and $\preceq_2$ on disjoint sets M_1 and M_2, we define a linear order $\preceq_1 + \preceq_2$ on $M_1 \cup M_2$, the *(ordered) sum* of $\preceq_1$ and $\preceq_2$, as $\preceq_1 \cup \preceq_2 \cup (M_1 \times M_2)$.

If $\varphi(\bar{y})$ is a formula and $\psi(\bar{x}, z)$ is a formula with at least one free variable z, then $\varphi|_\psi(\bar{x}, \bar{y})$ is the *relativisation of φ to ψ*. We construct $\varphi|_\psi$ by replacing subformulae $\exists x\, \chi$ and $\forall x\, \chi$ by $\exists x\, (\psi(\bar{y}, x) \wedge \varphi|_\psi)$ and $\forall z\, (\psi(\bar{y}, x) \rightarrow \varphi|_\psi)$, respectively. Note that $\mathrm{qad}(\varphi|_\psi) = \mathrm{qad}(\varphi)$ if ψ is an existential formula; in particular, $(\psi(\bar{y}, x) \rightarrow \varphi|_\psi) \equiv (\neg\psi(\bar{y}, x) \vee \varphi|_\psi)$ where, in this case, $\neg\psi(\bar{y}, x)$ is equivalent to a universal formula.

We transfer graph theoretic notions from graphs to general structures via the notion of Gaifman graphs. The *Gaifman graph* $\mathfrak{G}(\mathfrak{A})$ of a structure $\mathfrak{A}$ is the simple undirected graph with vertex set A containing an edge between $x, y \in A$ iff $x \neq y$ and x and y occur together in a tuple in one of the relations of $\mathfrak{A}$. The *distance* $\mathrm{dist}_{\mathfrak{A}}(a, b)$ between elements a, b of $\mathfrak{A}$ is their distance in $\mathfrak{G}(\mathfrak{A})$, i.e. the length of a shortest path between a and b in $\mathfrak{G}(\mathfrak{A})$. Similarly, notions such as *connectivity* and (connected) *components* of $\mathfrak{A}$ are defined. Note that the edge relation of the Gaifman graph is definable by an existential formula $\varphi_E(x, y)$, and this can be used to obtain, for every $\ell \geq 0$, an existential formula $\mathrm{dist}_{\leq\ell}(x, y)$ such that $\mathfrak{A} \models \mathrm{dist}_{\leq\ell}(a, b)$ iff $\mathrm{dist}_{\mathfrak{A}}(a, b) \leq \ell$.

Encoding information about elements in extended signatures In our proofs we will repeatedly remove single elements r from structures $\mathfrak{A}$ and encode information about the relations between r and the remaining elements into an expansion $\mathfrak{A}^{[r]}$ of the structure $\mathfrak{A} \setminus r$ (which is the substructure of $\mathfrak{A}$ induced on the elements different from r). We do this in such a way that the q-type of $\mathfrak{A}$ is determined by the q-type of $\mathfrak{A}^{[r]}$ together with what we call the atomic type of r in $\mathfrak{A}$.

The *atomic type* $\alpha(\mathfrak{A}, a)$ of an element a of a σ-structure $\mathfrak{A}$ is the set of all $R \in \sigma$ such that $(a, \ldots, a) \in R^{\mathfrak{A}}$ (where the tuple $(a, \ldots, a)$ has length $\mathrm{ar}(R)$). If no confusion seems likely we omit $\mathfrak{A}$ and just write $\alpha(a)$. Thus an atomic type is a subset of σ, and we identify $\alpha \subseteq \sigma$ with the $\mathrm{FO}[\sigma]$-sentence $\alpha(x) := \bigwedge_{R\in\alpha} R(x) \wedge \bigwedge_{R\in\sigma\setminus\alpha} \neg R(x)$.

Since we will often need the atomic type of the $\leq$-minimal element of a structure, we denote by $\alpha_{\mathfrak{A}}$ the type $\alpha(r, \mathfrak{A})$ if $\mathfrak{A}$ is an ordered structure with minimal element r.

To encode the relations between the element which is removed and the remaining elements, we define a signature $\tilde{\sigma}$ which contains, for each $R \in \sigma$ and each nonempty $I \subseteq [1, \mathrm{ar}(R)]$, a relation symbol R_I of arity $|I|$. Given a structure $\mathfrak{A} = (A, (R^{\mathfrak{A}})_{R\in\sigma})$ and an element $r \in A$ we now obtain a $\tilde{\sigma}$-structure $\mathfrak{A}^{[r]} = (A, (R_I^{\mathfrak{A}^{[r]}})_{R_I\in\tilde{\sigma}})$ by setting

$$R_I^{\mathfrak{A}^{[r]}} := \{(a_i)_{i\in I} \mid (a_1, \ldots, a_{\mathrm{ar}(R)}) \in R^{\mathfrak{A}} \text{ and } a_i = r \text{ for } i \notin I\}.$$

Note that $R^{\mathfrak{A}} = R^{\mathfrak{A}^{[r]}}_{[1,\mathrm{ar}(R)]}$, so up to a renaming of relation symbols, $\mathfrak{A}^{[r]}$ is an expansion of $\mathfrak{A} \setminus r$.

The (L, q)-type of $\mathfrak{A}$ is determined by $\alpha(r)$ and the (L, q)-type of $\mathfrak{A}^{[r]}$:

Lemma 1. *Let* $\mathrm{L} \in \{\mathrm{FO}, \mathrm{MSO}\}$ *and* $q \in \mathbb{N}^+$. *Let* $\mathfrak{A}$ *and* $\mathfrak{B}$ *be structures,* $r \in A$ *and* $s \in B$. *If*

$$\alpha(\mathfrak{A}, r) = \alpha(\mathfrak{B}, s) \quad \textit{and} \quad \mathrm{tp}_{\mathrm{L},q}(\mathfrak{A}^{[r]}) = \mathrm{tp}_{\mathrm{L},q}(\mathfrak{B}^{[s]}),$$

then also

$$\mathrm{tp}_{\mathrm{L},q}(\mathfrak{A}) = \mathrm{tp}_{\mathrm{L},q}(\mathfrak{B}).$$

Proof. The same argument works for $\mathrm{L} = \mathrm{FO}$ and $\mathrm{L} = \mathrm{MSO}$. Duplicator has a winning strategy $\mathcal{S}$ in the q-round Ehrenfeucht-Fraïssé game for L on $\mathfrak{A}^{[r]}$ and $\mathfrak{B}^{[s]}$. Note that the strategy $\mathcal{S}$ is, in particular, a winning strategy on $\mathfrak{A} \setminus r$ and $\mathfrak{B} \setminus s$, because $\mathfrak{A}^{[r]}$ and $\mathfrak{B}^{[s]}$ are expansions of these structures. Duplicator can win the q-round EF-game on $\mathfrak{A}$ and $\mathfrak{B}$ if she plays according to $\mathcal{S}$ on $\mathfrak{A} \setminus r$ and $\mathfrak{B} \setminus s$, and if she responds to r with s and vice versa.

We have to argue that this strategy preserves relations between the played elements. For relations not involving the removed elements r and s, this is true because S is a winning strategy for the q-round game on $\mathfrak{A} \setminus r$ and $\mathfrak{B} \setminus s$. Relations involving only the minimal elements are preserved because $\alpha(\mathfrak{A}, r) = \alpha(\mathfrak{B}, s)$. Relations involving the minimal elements and other elements are preserved, because they are encoded in the relations R_I of the extended signature $\tilde{\sigma}$, and these are preserved by $\mathcal{S}$. □

The following lemma is easy to prove following these definitions:

Lemma 2. *Let* $\mathrm{L} \in \{\mathrm{FO}, \mathrm{FO{+}MOD}\}$. *For every* $\mathrm{L}[\tilde{\sigma}]$*-sentence* φ *there is an* $\mathrm{L}[\sigma]$*-formula* $\mathcal{I}(\varphi)(z)$ *of the same quantifier rank and quantifier alternation depth such that*

$$\mathfrak{A} \models \mathcal{I}(\varphi)(r) \quad \textit{iff} \quad \mathfrak{A}^{[r]} \models \varphi,$$

for all σ*-structures* $\mathfrak{A}$ *and* $r \in A$.

Proof. The proof uses a standard interpretation argument. It suffices to provide quantifier-free formulae with a parameter z which define the universe and the relations of $\mathfrak{A}^{[r]}$ in $\mathfrak{A}$, provided that z is interpreted by the element r. The universe is defined by the formula $x \neq z$. Let $R_I \in \tilde{\sigma}$. If, for each $i \leq \mathrm{ar}(R)$, we let

$$y_i := \begin{cases} x_j & \text{if } i = i_j \in I \\ z & \text{if } i \notin I \end{cases}$$

then $R(y_1, \ldots, y_{\mathrm{ar}(R)})$ is a formula with free variables $z, x_1, \ldots, x_{|I|}$ which defines $R_I^{\mathfrak{A}^{[r]}}$ in $(\mathfrak{A}, r)$. □

Tree-depth The following inductive definition is one of several equivalent ways to define the *tree-depth* $\mathrm{td}(G)$ of a graph (see [16] for a reference on tree-depth):

$$\mathrm{td}(G) \;:=\; \begin{cases} 1 & \text{if } |V(G)| = 1 \\ 1 + \min_{r \in V(G)} \mathrm{td}(G \setminus r) & \text{if } G \text{ is connected and } |V(G)| > 1 \\ \max_{i \in [n]} \mathrm{td}(K_i) & \text{if } G \text{ has components } K_1, \dots, K_n. \end{cases}$$

As usual, the tree-depth $\mathrm{td}(\mathfrak{A})$ of a relational structure $\mathfrak{A}$ is defined by $\mathrm{td}(\mathfrak{A}) := \mathrm{td}(\mathfrak{G}(\mathfrak{A}))$. We let

$$\mathrm{Fin}_\sigma^{\mathrm{conn}} \;:=\; \{\mathfrak{A} \in \mathrm{Fin}_\sigma \mid \mathfrak{A} \text{ is connected}\}$$

and for each $d \in \mathbb{N}^+$, we let

$$\mathrm{Fin}_{\sigma,d} := \{\mathfrak{A} \in \mathrm{Fin}_\sigma \mid \mathrm{td}(\mathfrak{A}) \leq d\},$$
$$\mathrm{Fin}_{\sigma,d}^{\mathrm{conn}} := \{\mathfrak{A} \in \mathrm{Fin}_\sigma^{\mathrm{conn}} \mid \mathrm{td}(\mathfrak{A}) \leq d\}.$$

As an immediate consequence of the above definition of tree-depth, each $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}^{\mathrm{conn}}$ with $d > 1$ contains an element r with $\mathrm{td}(\mathfrak{A} \setminus r) \leq \mathrm{td}(\mathfrak{A}) - 1$. We call these vertices *tree-depth roots* and denote the set of all such vertices by $\mathrm{roots}(\mathfrak{A})$. By a result of [2], the size of $\mathrm{roots}(\mathfrak{A})$ is bounded by a function of d (independent of the size of $\mathfrak{A}$):

Lemma 3 ([2, Lem. 7])**.** *There is a function $f : \mathbb{N}^+ \to \mathbb{N}^+$ such that $|\mathrm{roots}(G)| \leq f(\mathrm{td}(G))$ for each connected graph G.*

Note that the definition of $\mathrm{roots}(G)$ in [2] is slightly different from ours, but the two definitions are easily seen to be equivalent.

A graph of tree-depth at most d can not contain a path of length 2^d (cf. [16, 6.2]). Therefore $\mathrm{dist}_{\mathfrak{A}}(a, b) < 2^d$ for all elements a and b in the same connected component of a structure $\mathfrak{A}$ of tree-depth at most d, and the formula $\mathrm{reach}_d(x, y) := \mathrm{dist}_{\leq 2^d}(x, y)$ defines the reachability relation in these structures:

$$\mathfrak{A} \models \mathrm{reach}_d[a, b] \quad \text{iff} \quad a \text{ and } b \text{ belong to the same component of } \mathfrak{A}.$$

This (existential) formula allows us to relativise a formula $\varphi(x)$ to the connected component of x:

$$\mathfrak{A} \models \varphi|_{\mathrm{reach}_d(x,z)}[a] \quad \text{iff} \quad \mathfrak{K} \models \varphi[a],$$

where $\mathfrak{K}$ is (the substructure of $\mathfrak{A}$ induced on) the connected component of a in $\mathfrak{A}$. Since reach_d is existential, we have $\mathrm{qad}(\varphi|_{\mathrm{reach}_d(x,z)}) = \mathrm{qad}(\varphi)$.

Using these observations and the inductive definition of tree-depth, it is easy to write down an $\mathrm{FO}[\sigma]$-sentence that defines $\mathrm{Fin}_{\sigma,d}$ on the class of all finite σ-structures. While this naïve approach leads to a formula whose quantifier alternation depth grows linearly with d, it is also possible to construct a *universal* sentence $\mathrm{td}_{\leq d}$ defining $\mathrm{Fin}_{\sigma,d}$ as a subclass of Fin_σ, cf. [16, Section 6.10] for details. Using this sentence, we construct a sentence that defines the set $\mathrm{roots}(\mathfrak{A})$ for each $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}^{\mathrm{conn}}$ with $d > 1$. To this end, we let $\mathrm{roots}_d(x) \;:=\; \bigvee_{c \leq d-1} \big(\mathrm{td}_{>c} \;\wedge\; \mathrm{td}_{\leq c}|_{(x \neq z)}(x)\big)$.

3 Order-invariant first-order logic

It is well-known that order-invariance is undecidable on the class Fin_σ of all finite σ-structures, i.e. there is no algorithm which decides for a given $\mathrm{FO}[\sigma^{\leq}]$-sentence if it is order-invariant on Fin_σ. This leads to the question if the expressive power of order-invariant sentences on a class $\mathcal{C}$ can be captured by a logic with a decidable syntax. An answer to this question in the case of the class Fin_σ seems out of reach. We consider the question in the case of bounded tree-depth structures, i.e. $\mathcal{C} = \mathrm{Fin}_{\sigma,d}$ for some $d \in \mathbb{N}^+$. More concretely, our aim is a proof of the following theorem:

Theorem 4. *For every $d \in \mathbb{N}^+$, every signature σ, and each sentence φ of $<$-inv-$\mathrm{FO}[\sigma]$, there is an $\mathrm{FO}[\sigma]$-sentence ψ which is equivalent to φ on $\mathrm{Fin}_{\sigma,d}$ and which has size $\|\psi\| \in d\text{-}\mathrm{EXP}(\mathrm{qr}(\varphi))$ and quantifier-alternation depth $\mathrm{qad}(\psi) \leq 3d$.*

The proof of Theorem 4 will be presented in Section 3.2 below. Before that, we want to motivate Theorem 4 by showing that the undecidability of order-invariance holds even for structures of tree-depth 2.

3.1 Undecidability of order-invariance on structures of tree-depth 2

As mentioned by [18], order-invariance on Fin_σ is decidable if the signature σ contains only unary relation symbols. An ordered σ-structure in which the unary relations partition the universe can be regarded as a word. An $\mathrm{FO}[\sigma^{\leq}]$-sentence φ then defines a language L_φ. The sentence φ is order-invariant iff the syntactic monoid of L_φ is commutative, which is decidable. This argument can be extended to general σ-structures and to structures of tree-depth 1 over arbitrary signatures.

Hence, order-invariance is decidable on $\mathrm{Fin}_{\sigma,d}$ if $d = 1$. The next theorem shows that it becomes undecidable for $d \geq 2$.

Theorem 5. *There is a signature σ such that order-invariance is undecidable on $\mathrm{Fin}_{\sigma,2}$.*

The proof of Theorem 5 uses a reduction from the undecidable halting problem for *counter machines* (cf. [15]) with two counters which store natural numbers. A counter machine executes a *program*, i.e. a finite sequence of the following *instructions*:

inc(i) increment counter i, proceed with next instruction.

dec(i, j_0, j_1) if counter i is not zero: decrement counter i, proceed with j_1-th instruction otherwise: proceed with instruction j_0.

halt stop the execution.

The configuration of the machine at any execution step is fully described by a triple (n_1, n_2, j), where $n_1, n_2 \geq 0$ are natural numbers stored in the counters and $j \geq 1$ is the number of the next instruction to be executed. Without loss of generality, we assume that the last instruction of a program is always the HALT instruction and that this instruction occurs nowhere else in the program. Hence we say that a program *halts* (on empty input) if it ever reaches its last instruction when run from the initial configuration $(0, 0, 1)$.

Proof of Theorem 5. We say that a sentence $\varphi \in \mathrm{FO}[\sigma^{\leq}]$ is *d-satisfiable* if it has a model $(\mathfrak{A}, \leq^{\mathfrak{A}})$ where $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$. The folklore proof which shows that order-invariance on Fin_σ is undecidable uses

a many-one reduction from the undecidable finite satisfiability problem to order-invariance. The same kind of argument proves that *d-satisfiability* (i.e. the problem which asks if a given sentence $\varphi \in \text{FO}[\sigma^{\leq}]$ is d-satisfiable) many-one reduces to order-invariance on $\text{Fin}_{\tilde{\sigma},d}$, where $\tilde{\sigma} := \sigma \cup \{P\}$ for a unary relation symbol $P \notin \sigma$. This follows from the fact that $\varphi \in \text{FO}[\sigma^{\leq}]$ is d-satisfiable if, and only if, the $\text{FO}[\tilde{\sigma}^{\leq}]$-sentence $\varphi \wedge \exists x \forall y\ (x{\leq}y\ \wedge P(x))$ is *not* order-invariant on $\text{Fin}_{\tilde{\sigma},d}$.

Hence, to complete the proof of our theorem, it suffices to show that the 2-satisfiability problem is undecidable for some signature σ to be fixed below. To this end we reduce the halting problem for counter machines to 2-satisfiability. Let $P = I_1 \cdots I_\ell$ be a program. We construct an $\text{FO}[\sigma^{\leq}]$-sentence φ which is $\text{Fin}_{\sigma,2}$-satisfiable iff P halts. First we fix an encoding of configurations of P by words over a finite alphabet Σ. It would be natural to do this by encoding the counter values in unary using different symbols; say, $(2, 3, 1)$ would become $\mathbf{11}\,\mathbf{222}\,1$. We change this representation slightly: a configuration (n_1, n_2, j) of P is encoded by a word

$$\text{enc}(n_1, n_2, j) \ := \ (\mathbf{1}_L \mathbf{1}_R)^{n_1} \, (\mathbf{2}_L \mathbf{2}_R)^{n_2} \, j$$

over the alphabet $\Sigma := \{\mathbf{1}_L, \mathbf{2}_L, \mathbf{1}_R, \mathbf{2}_R, 1, \ldots, \ell\}$. [1]

Let $\sigma := \{E\} \cup \tau$ where E is a binary relation symbol and $\tau := \{P_a \mid a \in \Sigma\}$, where the P_a are unary relation symbols. The σ-structures that we consider are Σ-*coloured graphs*, i.e. σ-structures where E is the edge relation of a simple undirected graph and where the unary predicates are a vertex colouring (i.e. a partition of the vertex set). If a vertex of such a graph belongs to a relation P_a, we say that it is *a-coloured*. The class of Σ-coloured graphs is obviously FO-definable on Fin_σ.

As usual, we identify each non-empty word over the alphabet Σ with an ordered τ-structure which, in turn, we regard as an ordered Σ-coloured graph with no edges. We refer to vertices which are coloured by $1, \ldots, \ell$ as *instruction vertices*. If our program P halts after at most h computation steps then, with respect to our encoding, there exists a unique word w_P which encodes the *run* of P, i.e. the finite sequence of configurations at time steps $1, \ldots, h$. We want to define a class of ordered Σ-coloured graphs of maximum degree 1 obtained from the edge-less graph w_P by adding edges between its vertices. These graphs will be called *matching extensions* of w_P, since their edge relations will be unions of matchings (i.e. edge relations of graphs where each vertex is incident to exactly one edge). Consider any word $w = \text{enc}(C_1) \cdots \text{enc}(C_k)$ which encodes a sequence of representations. We phrase the description of the execution of the counter machine program P given in the definition of counter machines above somewhat more formally as conditions under which the sequence $C_1, \ldots, C_k$ is a run of P (i.e. $w = w_P$). At the same time, we rephrase them as statements about the ordered Σ-coloured graph w in a way that will be suitable for the definition of our sentence φ.

1. $C_1 = (0, 0, 1)$ and C_k is a halting configuration, i.e. $C_k = (n_1, n_2, \ell)$ for some $n_1, n_2 \geq 0$.

   With our encoding, this is equivalent to the first vertex of w being 1-coloured and the last vertex being ℓ-coloured. (Recall that the machine starts with both counters being 0.)

2. For each $i \in [k-1]$ and $C_i = (n_1, n_2, j)$ one of the following statements is true:

   (a) $I_j = \text{INC}(1)$ and $C_{i+1} = (n_1 + 1, n_2, j + 1)$.

[1] This alphabet depends on the length of the given program P, but the proof can be modified easily to make the alphabet Σ, and therefore the signature σ, independent of P without increasing the tree-depth of the structures involved.

This holds iff we can add edges to w so that all $\mathbf{1}_L$-coloured vertices in $\mathrm{enc}(C_i)$ are matched with all but one of the $\mathbf{1}_R$-coloured vertices in $\mathrm{enc}(C_{i+1})$, and the $\mathbf{2}_L$-coloured vertices in $\mathrm{enc}(C_i)$ are matched with the $\mathbf{2}_R$-coloured vertices in $\mathrm{enc}(C_{i+1})$, and the unique instruction vertices in $\mathrm{enc}(C_i)$ and $\mathrm{enc}(C_{i+1})$ have the same colour.

(b) $I_j = \textsc{dec}(1, j_0, j_1)$ and either $n_1 = 0$ and $C_{i+1} = (n_1, n_2, j_0)$, or $n_1 \geq 1$ and $C_{i+1} = (n_1 - 1, n_2, j_1)$.

Equivalently, either one of the following statements is true:

- There exists no $\mathbf{1}_L$-coloured vertex in $\mathrm{enc}(C_i)$ and no $\mathbf{1}_L$-coloured vertex in $\mathrm{enc}(C_{i+1})$. Furthermore, the $\mathbf{2}_L$-coloured vertices in $\mathrm{enc}(C_i)$ can be matched with the $\mathbf{2}_R$-coloured vertices in $\mathrm{enc}(C_{i+1})$. The unique instruction vertex in $\mathrm{enc}(C_{i+1})$ is j_0-coloured.
- There is at least one $\mathbf{1}_L$-coloured vertex in $\mathrm{enc}(C_i)$. Furthermore, the $\mathbf{1}_R$-coloured vertices in $\mathrm{enc}(C_{i+1})$ can be matched with all but one of the $\mathbf{1}_L$-coloured vertices in $\mathrm{enc}(C_i)$, and the $\mathbf{2}_L$-coloured vertices in $\mathrm{enc}(C_i)$ can be matched with the $\mathbf{2}_R$-coloured vertices in $\mathrm{enc}(C_{i+1})$. The unique instruction vertex in $\mathrm{enc}(C_{i+1})$ is j_1-coloured.

(c),(d) Analogous statements to (a), (b) for the case where I_j operates on counter 2.

Now, a *matching extension* of w_P is an ordered graph obtained from w_P by adding, for each pair of subsequent configurations, exactly the edges of a matching witnessing that w_P satisfies the conditions (a), (b), (c), and (d). Observe that each vertex of a matching extension is contained in at most one matching. Hence, any matching extension has maximum degree 1. Using our description above, it is easy to write down a first-order sentence φ defining the class of all matching extensions of w_P. This class is non-empty iff P halts. Hence φ is 2-satisfiable iff P halts. □

3.2 From order-invariant FO[$\sigma^{\leq}$]-formulae to FO[σ]-formulae

We prove Theorem 4. The key insight here is that for every quantifier rank q and every structure $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$ there exists a class of canonical linear orders $\preceq_q$ for which the FO_q-type of $(\mathfrak{A}, \preceq_q)$ is already FO-definable in $\mathfrak{A}$. In particular, $\mathrm{tp}_q(\mathfrak{A}, \preceq_q)$ only depends on $\mathfrak{A}$, even though there may be more than one such order on $\mathfrak{A}$.

We call these canonical orders *q-orders*. After defining them formally we will thus prove the following two facts about them:

1. Expansions by q-orders are indistinguishable in FO_q, i.e. $(\mathfrak{A}, \preceq_1) \equiv_q (\mathfrak{A}, \preceq_2)$ for all finite structures $\mathfrak{A}$, provided both $\preceq_1$ and $\preceq_2$ are q-orders (cf. Lemma 6).

2. If the tree-depth of structures is bounded, then the q-type $\mathrm{tp}_q(\mathfrak{A}, \preceq_q)$ of an expansion of $\mathfrak{A}$ by a q-order is definable in FO (Lemmas 10 and 13). The proof of Theorem 4 easily follows from this.

The definition of q-orders With an eye towards Section 4, the notion of q-orders will be defined more generally for logics $\mathrm{L} \in \{\mathrm{FO}, \mathrm{MSO}\}$. We fix arbitrary orders $\preceq_{\mathrm{L},q}$ on the set of (L, q)-types over the signature $\sigma^{\leq}$, and $\preceq_{\mathrm{atomic}}$ on the set of atomic σ-types. For simplicity we write $a \preceq_{\mathrm{atomic}} b$ for $\alpha(a) \preceq_{\mathrm{atomic}} \alpha(b)$.

To obtain a q-order $\preceq$ on a connected structure $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$, we pick a root r of $\mathfrak{A}$ which has $\preceq_{\mathrm{atomic}}$-minimal atomic type among all roots and for which the type of q-ordered expansions of $\mathfrak{A}^{[r]}$ is $\preceq_{\mathrm{L},q}$-minimal among all $\preceq_{\mathrm{atomic}}$-minimal roots. We place this r in front of the order $\preceq$ and order the remaining elements according to a (recursively obtained) q-order on $\mathfrak{A}^{[r]}$. On structures with more than one component, we q-order the components individually and take the sum of their orders, following the $\preceq_{\mathrm{L},q}$-order of the components:

Definition 1 ((L,q)-order)**.** *An* (L,q)-order *on a* σ*-structure* $\mathfrak{A}$ *is an order* $\preceq$ *which satisfies the following conditions:*

(1) *If* $\mathfrak{A}$ *is connected we denote by* $r \in A$ *its* $\preceq$*-minimal element. Then either* $|A| = 1$*, or* $|A| > 1$ *and the following holds:*

  (1) r *is a* $\preceq_{atomic}$*-minimal root of* $\mathfrak{A}$*, i.e.* $r \in \mathrm{roots}(\mathfrak{A})$ *and* $r \preceq_{atomic} r'$ *for all* $r' \in \mathrm{roots}(\mathfrak{A})$*.*

  (2) *The* (L,q)*-type of* q*-ordered expansions of* $\mathfrak{A}^{[r]}$ *is minimal:*

$$\mathrm{tp}_q(\mathfrak{A}^{[r]}, \preceq) \preceq_{\mathrm{L},q} \mathrm{tp}_q(\mathfrak{A}^{[r']}, \preceq')$$

  for every $r' \in \mathrm{roots}(\mathfrak{A})$ *with* $\alpha(r') = \alpha(r)$ *and every* q*-order* $\preceq'$ *on* $\mathfrak{A}^{[r']}$*.*

  (3) $\preceq|_{A\setminus r}$ *is an* (L,q)*-order on* $\mathfrak{A}^{[r]}$*.*

(2) *If* $\mathfrak{A}$ *is not connected, we denote its components by* $\mathfrak{A}_1, \dots, \mathfrak{A}_\ell$ *and set* $\preceq_i := \preceq|_{A_i}$*. Then* $\preceq$ *is a* q*-order if*

  (1) *each* $\preceq_i$ *is a* q*-order of* $\mathfrak{A}_i$*, and*

  (2) *after suitably permuting the components,*

$$\preceq = \preceq_1 + \cdots + \preceq_\ell \quad \text{and} \quad \mathrm{tp}_q(\mathfrak{A}_i, \preceq_i) \preceq_{\mathrm{L},q} \mathrm{tp}_q(\mathfrak{A}_j, \preceq_j) \text{ for } i \le j.$$

The $\preceq$*-minimal element of a* q*-order* $\preceq$ *will be denoted by* $r_\preceq$*.*

It is plain from the definition above that each structure can be q-ordered. Next we want to show that all q-ordered expansions $(\mathfrak{A}, \preceq)$ of a given structure $\mathfrak{A}$ have the same q-type, and that the q-type of $(\mathfrak{A}^{[r_\preceq]}, \preceq)$ is also the same for all q-orders $\preceq$ of $\mathfrak{A}$.

Lemma 6. *Let* $\mathrm{L} \in \{\mathrm{FO}, \mathrm{MSO}\}$*,* $q \in \mathbb{N}^+$*. For all* (L,q)*-orders* $\preceq, \preceq'$ *of a structure* $\mathfrak{A}$*, we have*

$$(\mathfrak{A}, \preceq) \equiv_q^{\mathrm{L}} (\mathfrak{A}, \preceq').$$

If $\mathfrak{A}$ *is connected and* $\mathrm{td}(A) > 1$*, then also* $(\mathfrak{A}^{[r_\preceq]}, \preceq) \equiv_q^{\mathrm{L}} (\mathfrak{A}^{[r_{\preceq'}]}, \preceq')$*.*

For the proof, we will need the following composition lemma for ordered sums, cf. [14] for a proof.

Lemma 7 (Composition Lemma)**.** *Let* $\mathrm{L} \in \{\mathrm{FO}, \mathrm{MSO}\}$*,* $q \in \mathbb{N}$ *and let* σ *be a relational signature. Let* $(\mathfrak{A}_1, \preceq^{\mathfrak{A}_1})$*,*$(\mathfrak{A}_2, \preceq^{\mathfrak{A}_2})$*,*$(\mathfrak{B}_1, \preceq^{\mathfrak{B}_1})$*,*$(\mathfrak{B}_2, \preceq^{\mathfrak{B}_2})$ *be ordered* σ*-structures. If*

$$(\mathfrak{A}_1, \preceq^{\mathfrak{A}_1}) \equiv_q^{\mathrm{L}} (\mathfrak{A}_2, \preceq^{\mathfrak{A}_2}) \quad \text{and} \quad (\mathfrak{B}_1, \preceq^{\mathfrak{B}_1}) \equiv_q^{\mathrm{L}} (\mathfrak{B}_2, \preceq^{\mathfrak{B}_2}),$$

then

$$(\mathfrak{A}_1 \sqcup \mathfrak{B}_1, \preceq^{\mathfrak{A}_1} + \preceq^{\mathfrak{B}_1}) \equiv_q^{\mathrm{L}} (\mathfrak{A}_2 \sqcup \mathfrak{B}_2, \preceq^{\mathfrak{A}_2} + \preceq^{\mathfrak{B}_2}).$$

Proof of Lemma 6. The proof proceeds on the size of A. If $|A| = 1$ then $\preceq = \preceq'$ and there is nothing to prove.

Let $|A| > 1$ and suppose first that $\mathfrak{A}$ is connected. By Definition 1, $\alpha(r_\preceq) = \alpha(r_{\preceq'})$ and

$$\mathrm{tp}_q(\mathfrak{A}^{[r_\preceq]}, \preceq) \preceq_{\mathrm{L},q} \mathrm{tp}_q(\mathfrak{A}^{[r_{\preceq'}]}, \preceq').$$

By symmetry also

$$\mathrm{tp}_q(\mathfrak{A}^{[r_{\preceq'}]}, \preceq') \preceq_{\mathrm{L},q} \mathrm{tp}_q(\mathfrak{A}^{[r_\preceq]}, \preceq),$$

so $\mathrm{tp}_q(\mathfrak{A}^{[r_\preceq]}, \preceq) = \mathrm{tp}_q(\mathfrak{A}^{[r_{\preceq'}]}, \preceq')$ and, by Lemma 1, $(\mathfrak{A}, \preceq) \equiv_q (\mathfrak{A}, \preceq')$.

Now consider the case where $\mathfrak{A}$ is not connected, and let $\mathfrak{K}_1, \dots, \mathfrak{K}_\ell$ be the components of $\mathfrak{A}$. By the definition of q-orders each $\mathfrak{K}_i$ is q-ordered, so

$$(\mathfrak{K}_i, \preceq|_{K_i}) \equiv^{\mathrm{L}}_q (\mathfrak{K}_i, \preceq'|_{K_i})$$

for $i = 1, \dots, \ell$ by what we have just said. Considering the way that an (L, q)-order orders the components of a structure according to their (L, q)-types (Part 2 of Definition 1), we obtain that $(\mathfrak{A}, \preceq) \equiv^{\mathrm{L}}_q (\mathfrak{A}, \preceq')$ by repeatedly applying the Composition Lemma. □

By Lemma 6 it makes sense to speak of the *q-order type* of an unordered structure $\mathfrak{A}$ which we define as $\mathrm{tp}^{\leq}_q(\mathfrak{A}) := \mathrm{tp}_q(\mathfrak{A}, \preceq_q)$ If $\mathfrak{A}$ is connected and $\mathrm{td}(\mathfrak{A}) > 1$, we furthermore define its *q-order root type* as $\mathrm{rtp}^{\leq}_q(\mathfrak{A}) := \mathrm{tp}_q(\mathfrak{A}^{[r_{\preceq_q}]}, \preceq_q)$. In both cases $\preceq_q$ is some q-order on $\mathfrak{A}$ and well-definedness is guaranteed by the Lemma. Note that both these types are $\sigma^{\leq}$-types. Similarly, the atomic type $\alpha_\mathfrak{A} := \alpha(r_\leq)$ of the minimal element in a q-ordered expansion of $\mathfrak{A}$ is well-defined.

We set

$$\begin{aligned}
\mathcal{T}_{\mathrm{L},\sigma,q,d} &:= \{\mathrm{tp}^{\leq}_q(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{Fin}_{\sigma,d}\},\\
\mathcal{T}^{\mathrm{conn}}_{\mathrm{L},\sigma,q,d} &:= \{\mathrm{tp}^{\leq}_q(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{Fin}^{\mathrm{conn}}_{\sigma,d}\}, \text{ and}\\
\mathcal{T}_{\mathrm{L},\sigma,q} &:= \bigcup_{d \in \mathbb{N}^+} \mathcal{T}_{\mathrm{L},q,\sigma,d}.
\end{aligned}$$

We say that a sentence $\varphi_\tau \in \mathrm{L}[\sigma]$ *defines* τ *on* $\mathrm{Fin}_{\sigma,d}$ (and that τ is L-definable) if for each $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$, we have

$$\mathfrak{A} \models \varphi_\tau \quad \text{iff} \quad \mathrm{tp}^{\leq}_q(\mathfrak{A}) = \tau.$$

Note that the sentence φ_τ must not contain the relation $\leq$.

By Lemma 1 the atomic type of $r_\preceq$ and the q-type of $\mathfrak{A}^{[r_\preceq]}$ determine the q-type of $\mathfrak{A}$, and $\mathrm{td}(\mathfrak{A}^{[r_\preceq]}) = \mathrm{td}(\mathfrak{A}) - 1$, for connected structures $\mathfrak{A}$ and q-orders $\preceq$. Since the number of atomic $\tilde{\sigma}$-types is $2^{|\tilde{\sigma}|}$, we obtain the following bound on the size of $\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}$:

Corollary 8. *Let $q, d \in \mathbb{N}^+$. Then $|\mathcal{T}^{conn}_{\sigma,q,d}| \leq 2^{|\tilde{\sigma}|} \cdot |\mathcal{T}_{\tilde{\sigma},q,d-1}|$.*

3.3 Handling connected structures

The proof of our main theorem is broken down into two steps. In the first step, we show how to lift the definability of q-types of q-ordered structures from structures of tree-depth $d-1$ to connected structures of tree-depth d.

Again we invoke Lemma 1 and Lemma 6 to show that q-order types can be broken down into atomic types of roots and q-order root types:

Corollary 9. *Let $d > 1$ and let $\tau \in \mathcal{T}^{conn}_{\sigma,q,d}$. Let*

$$R_\tau \;:=\; \{(\alpha_{\mathfrak{A}}, \mathrm{rtp}^{\leq}_q(\mathfrak{A})) \mid \mathfrak{A} \in \mathrm{Fin}^{conn}_{\sigma,d}, \mathrm{td}(\mathfrak{A}) > 1, \text{ and } \mathrm{tp}^{\leq}_q(\mathfrak{A}) = \tau\}.$$

Then for each $\mathfrak{B} \in \mathrm{Fin}^{conn}_{\sigma,d}$, we have $\mathrm{tp}^{\leq}_q(\mathfrak{B}) = \tau$ iff $(\alpha_{\mathfrak{B}}, \mathrm{rtp}^{\leq}_q(\mathfrak{B})) \in R_\tau$.

Proof. The "only-if"-part of the claim is obvious. Regarding the "if"-part, if

$$(\alpha_{\mathfrak{B}}, \mathrm{rtp}^{\leq}_q(\mathfrak{B})) = (\alpha_{\mathfrak{A}}, \mathrm{rtp}^{\leq}_q(\mathfrak{A}))$$

for some $\mathfrak{A}$ with $\mathrm{tp}^{\leq}_q(\mathfrak{A}) = \tau$, then Lemma 6 and the definitions of $\mathrm{tp}^{\leq}_q, \mathrm{rtp}^{\leq}_q$ imply that $\mathrm{tp}^{\leq}_q(\mathfrak{B}) = \tau$. □

Lemma 10. *Let $q, d \in \mathbb{N}^+$ with $d > 1$. Let $(\mathrm{L}_1, \mathrm{L}_2)$ be one of $(\mathrm{FO}, \mathrm{FO})$ or $(\mathrm{MSO}, \mathrm{FO{+}MOD})$. If each (L_1, q)-type $\theta \in \mathcal{T}_{\tilde{\sigma},q,d-1}$ is $\mathrm{L}_2[\tilde{\sigma}]$-definable on $\mathrm{Fin}_{\tilde{\sigma},d-1}$ by a sentence $\psi_{\theta,d-1}$, then each (L_1, q)-type $\tau \in \mathcal{T}^{conn}_{\sigma,q,d}$ is $\mathrm{L}_2[\sigma]$-definable on $\mathrm{Fin}^{conn}_{\sigma,d}$ by a sentence $\varphi^{conn}_{\tau,d}$. Moreover, defining*

$$\Psi := \{\psi_{\theta,d-1} \mid \theta \in \mathcal{T}_{\tilde{\sigma},q,d-1}\} \quad \text{and} \quad \Phi := \{\varphi^{conn}_{\tau,d} \mid \tau \in \mathcal{T}_{\sigma,q,d}\},$$

we have $\|\Phi\| \leq c \cdot \|\Psi\| \cdot |\mathcal{T}_{\tilde{\sigma},q,d-1}|^2$ and $\mathrm{qad}(\Psi) \leq \mathrm{qad}(\Phi) + 1$, for a constant c depending only on σ, d.

Proof. In the following, all q-types are $(\mathrm{L}_1, (\sigma^{\leq}), q)$-types. Let $\tau \in \mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}$ and let R_τ be as in Corollary 9. We show that, under the assumptions of our lemma, the class

$$\{\mathfrak{A} \in \mathrm{Fin}^{\mathrm{conn}}_{\sigma,d} \mid (\alpha_{\mathfrak{A}}, \mathrm{rtp}^{\leq}_q(\mathfrak{A})) \in R_\tau\}$$

is $\mathrm{L}_2[\sigma]$-definable by a sentence φ_τ on $\mathrm{Fin}^{\mathrm{conn}}_{\sigma,d}$. Taking care of connected structures of tree-depth 1 (i.e. singleton structures) we set $\varphi^{\mathrm{conn}}_{\tau,d} := (\mathrm{td}_{\leq 1} \wedge \hat{\varphi}_\tau) \vee (\mathrm{td}_{>1} \wedge \varphi_\tau)$, where $\hat{\varphi}_\tau$ defines τ on singleton structures.

For each atomic σ-type $\alpha \subseteq \sigma$, the following FO-sentence ξ_α expresses in a structure $\mathfrak{A} \in \mathrm{Fin}^{\mathrm{conn}}_{\sigma,d}$ that $\alpha_{\mathfrak{A}} = \alpha$:

$$\xi_\alpha \;:= \big(\exists x \,\big(\mathrm{roots}_d(x) \wedge \alpha(x)\big)\big) \wedge \left(\forall x \,\big(\mathrm{roots}_d(x) \to \bigvee_{\alpha \preceq_{\mathrm{atomic}} \alpha'} \alpha'(x)\big)\right).$$

For each type $\theta \in \mathcal{T}_{\tilde{\sigma},q,d-1}$ the following sentence is true in a σ-structure $\mathfrak{A}$ if, and only if, there is a root r of atomic type α for which $\mathfrak{A}^{[r]}$ has type θ, and θ is $\preceq_{\mathrm{L}_1,q}$-minimal among the types of $\mathfrak{A}^{[s]}$ for roots s of atomic type α:

$$\begin{aligned}\chi_{\alpha,\theta} \;&:= \forall x \left((\mathrm{roots}_d(x) \wedge \alpha(x)) \to \bigvee_{\theta \preceq_{\mathrm{L}_1,q} \theta'} \mathcal{I}(\psi_{\theta',d-1})(x)\right) \\ &\wedge \exists x \,\big(\mathrm{roots}_d(x) \wedge \alpha(x) \wedge \mathcal{I}(\psi_{\theta,d-1})(x)\big).\end{aligned}$$

Observe that $\mathrm{qad}(\chi_{\alpha,\theta}) \leq \mathrm{qad}(\Psi) + 1$.

Now we obtain the desired sentence by defining $\varphi_\tau := \bigvee_{(\alpha,\theta) \in R_\tau} (\xi_\alpha \wedge \chi_{\alpha,\theta})$.

Observe that, for some constant c depending only on σ, d, we have $\|\xi_\alpha\| \leq c$, $\|\chi_{\alpha,\theta}\| \leq c \cdot \|\Psi\| \cdot |\mathcal{T}_{\tilde{\sigma},q,d-1}|$, $|R_\tau| \leq c \cdot |\mathcal{T}_{\tilde{\sigma},q,d-1}|$, and $\|\varphi_\tau\| \leq c \cdot \|\Psi\| \cdot |\mathcal{T}_{\tilde{\sigma},q,d-1}|^2$. The claims about $\|\Phi\|$ and $\mathrm{qad}(\Phi)$ follow from the observations above. □

3.4 Handling disconnected structures

We proceed with the preparations for the second step in the proof of our main theorem, where we lift the definability of q-order types from connected structures of tree-depth $\leq d$ to disconnected structures of tree-depth $\leq d$.

For us, a *Boolean query* is an isomorphism-invariant map $f : \mathrm{Fin} \to \{0,1\}$, where Fin is the class of all finite structures (i.e. structures over arbitrary signatures). We will treat maps $f : \mathrm{Fin}_\sigma \to \{0,1\}$ as Boolean queries by assuming that $f(\mathfrak{A}) = 0$ if $\mathfrak{A}$ is not a σ-structure. The general definition for arbitrary signatures will be useful in in Section 5 below. We are interested in two kinds of queries. As usual, we identify each sentence φ with a Boolean query such that $\varphi(\mathfrak{A}) = 1$ iff $\mathfrak{A} \models \varphi$. Furthermore, we identify each q-order type τ with a query such that $\tau(\mathfrak{A}) = 1$ iff $\mathrm{tp}_q^{\leq}(\mathfrak{A}) = \tau$. For each structure $\mathfrak{A}$ and each Boolean query f, we let $n_f(\mathfrak{A})$ denote the number of components $\mathfrak{K}$ of $\mathfrak{A}$ such that $f(\mathfrak{K}) = 1$. For each ordered set $Q := \{f_1, \dots, f_\ell\}$ of Boolean queries, we let $\bar{n}_Q(\mathfrak{A}) := (n_{f_1}(\mathfrak{A}), \dots, n_{f_\ell}(\mathfrak{A}))$. For natural numbers $a, b, t \in \mathbb{N}^+$ we set

$$a \equiv_{\wedge t} b \quad \Leftrightarrow \quad (a = b \text{ or } a, b \geq t),$$

and we extend this relation to tuples $\bar{a}$ and $\bar{b}$ by saying $\bar{a} \equiv_{\wedge t} \bar{b}$ if, and only if, $a_i \equiv_{\wedge t} b_i$ for all components a_i and b_i.

We show that FO inherits its capability to count the types of components in q-ordered structures from its capability to distinguish linear orders of different length. The proof of the following lemma closely follows a step in the proof of [1, Thm. 5.5]. Observe that for all $\mathfrak{A}, \mathfrak{B} \in \mathrm{Fin}_{\sigma,d}$, $n_{\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}}(\mathfrak{A}) \equiv_{\wedge t} n_{\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}}(\mathfrak{B})$ iff $n_{\mathcal{T}_{\sigma,q}}(\mathfrak{A}) \equiv_{\wedge t} n_{\mathcal{T}_{\sigma,q}}(\mathfrak{B})$.

Lemma 11. *Let $d \geq 1$, $q \in \mathbb{N}^+$ and $t := 2^q + 1$. Then for all $\mathfrak{A}, \mathfrak{B} \in \mathrm{Fin}_{\sigma,d}$,*

$$n_{\mathcal{T}_{\sigma,q}}(\mathfrak{A}) \equiv_{\wedge t} n_{\mathcal{T}_{\sigma,q}}(\mathfrak{B}) \implies \mathrm{tp}_q^{\leq}(\mathfrak{A}) = \mathrm{tp}_q^{\leq}(\mathfrak{B}).$$

Proof. For each component $\mathfrak{K}$ of $\mathfrak{A}$, we let $\preceq^{\mathfrak{K}}$ be a q-order of $\mathfrak{K}$. By Part 2 of Definition 1, the q-orders on the components of $\mathfrak{A}$ can be extended to a q-order $\preceq^{\mathfrak{A}}$ on $\mathfrak{A}$ such that $\preceq^{\mathfrak{A}}|_{\mathfrak{K}} = \preceq^{\mathfrak{K}}$ for each component $\mathfrak{K}$ of $\mathfrak{A}$. We proceed analogously to obtain a q-order $\preceq^{\mathfrak{B}}$ on $\mathfrak{B}$. Let $\mathcal{T}_{\sigma,q} = \{\tau_1, \dots, \tau_\ell\}$, where $\ell := |\mathcal{T}_{\sigma,q}|$ and $\tau_i \preceq_q \tau_j$ iff $i < j$. We consider words over the alphabet $\mathcal{T}_{\sigma,q}$ as structures in the usual way, i.e. as ordered structures over a signature containing a unary relation symbol for each type. Consider the words $w_{\mathfrak{A}}, w_{\mathfrak{B}} \in \mathcal{T}^*_{\sigma,q}$ obtained from $(\mathfrak{A}, \preceq^{\mathfrak{A}})$ and $(\mathfrak{B}, \preceq^{\mathfrak{B}})$ by contracting each component $\mathfrak{K}$ to a single element that gets labelled by its q-type in the corresponding q-ordered structure. By this construction and by Part 2 of Definition 1, we know that

$$w_{\mathfrak{A}} = \tau_1^{n_{\tau_1}(\mathfrak{A})} \cdots \tau_\ell^{n_{\tau_\ell}(\mathfrak{A})} \quad \text{and} \quad w_{\mathfrak{B}} = \tau_1^{n_{\tau_1}(\mathfrak{B})} \cdots \tau_\ell^{n_{\tau_\ell}(\mathfrak{B})}.$$

Since $n_{\mathcal{T}_{\sigma,q}}(\mathfrak{A}) \equiv_{\wedge t} n_{\mathcal{T}_{\sigma,q}}(\mathfrak{B})$, for each $i \in [\ell]$, we have either $n_{\tau_i}(\mathfrak{A}) = n_{\tau_i}(\mathfrak{B})$ or $n_{\tau_i}(\mathfrak{A}), n_{\tau_i}(\mathfrak{B}) \geq t$. A folklore result (cf. [13, Ch. 3]) tells us that $w_{\mathfrak{A}} \equiv_q^{\mathrm{FO}} w_{\mathfrak{B}}$, i.e. Duplicator has a winning strategy in the q-round EF-game on the two word structures.

We show that $(\mathfrak{A}, \preceq^{\mathfrak{A}}) \equiv_q^{\mathrm{FO}} (\mathfrak{B}, \preceq^{\mathfrak{B}})$. To this end, consider the following winning strategy for Duplicator in the q-round EF-game on $(\mathfrak{A}, \preceq^{\mathfrak{A}})$ and $(\mathfrak{B}, \preceq^{\mathfrak{B}})$. She maintains a *virtual* q-round EF-game $w_{\mathfrak{A}}$ on $w_{\mathfrak{B}}$ between a *Virtual Spoiler* and a *Virtual Duplicator*. When, during the i-th round, Spoiler chooses an element v in some component $\mathfrak{K}$ of, say, $\mathfrak{A}$, she lets the Virtual Spoiler play the corresponding position in $w_{\mathfrak{A}}$ in the i-th round of the virtual game. The Virtual Duplicator answers

in $w_{\mathfrak{B}}$. Duplicator chooses a component $\mathfrak{K}'$ of $\mathfrak{B}$ for its reply according to the Virtual Duplicator's answer in $w_{\mathfrak{B}}$. The winning strategy on $w_{\mathfrak{A}}$ and $w_{\mathfrak{B}}$ ensures that $(\mathfrak{K}, \preceq^{\mathfrak{A}}) \equiv_q^{\mathrm{FO}} (\mathfrak{K}', \preceq^{\mathfrak{B}})$ and that all elements of $\mathfrak{K}$ and $\mathfrak{K}'$ have the same positions in $\preceq^{\mathfrak{A}}$ and $\preceq^{\mathfrak{B}}$ relative to the elements played in the previous rounds. Duplicator uses her winning strategy in the q-round game on the ordered components to determine the element of $\mathfrak{K}'$ that she uses as her answer to v. □

For a tuple $\bar{a}$ of natural numbers, denote by $[\bar{a}]_{\wedge t}$ the tuple obtained from it by replacing all entries $> t$ with t. Then the previous lemma implies that if $\mathrm{td}(\mathfrak{A}) \leq d$, then $[\bar{n}_{\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}}(\mathfrak{A})]_{\wedge(2^q+1)}$ determines $\mathrm{tp}^{\leq}(\mathfrak{A})$. Hence we obtain the following corollary:

Corollary 12. *Let $q, d \in \mathbb{N}^+$ and let $t := 2^q + 1$. For each $\varphi \in \mathrm{FO}[\sigma^{\leq}]$, let*

$$R_\varphi \;:=\; \{[\bar{n}_{\mathcal{T}^{conn}_{\sigma,q,d}}(\mathfrak{A})]_{\wedge t} \mid \mathfrak{A} \in \mathrm{Fin}_{\sigma,d}, \mathrm{tp}^{\leq}_q(\mathfrak{A}) \models \varphi\}.$$

Then for each $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$, we have

$$\mathrm{tp}^{\leq}_q(\mathfrak{A}) \models \varphi \quad \textit{if, and only if,} \quad [\bar{n}_{\mathcal{T}^{conn}_{\sigma,q,d}}]_{\wedge t} \in R_\tau.$$

Furthermore, $|\mathcal{T}_{\sigma,q,d}| \leq (t+1)^{|\mathcal{T}^{conn}_{\sigma,q,d}|}$.

The following lemma will be used in conjunction with the previous corollary to lift the definability of q-types from connected to disconnected structures.

Lemma 13. *Let $\mathrm{L} \in \{\mathrm{FO}, \mathrm{FO{+}MOD}\}$. For all $d, t \in \mathbb{N}^+$, every set of L-sentences Φ, and every set $R \subseteq [0,t]^{|\Phi|}$, there is an L-sentence ψ^Φ_R such that for each structure $\mathfrak{A}$ with $\mathrm{td}(\mathfrak{A}) \leq d$, we have*

$$\mathfrak{A} \models \psi^\Phi_R \iff [\bar{n}_\Phi(\mathfrak{A})]_{\wedge t} \in R.$$

Moreover, $\|\psi^\Phi_R\| \leq c \cdot |\Phi| \cdot \|\Phi\| \cdot |R| \cdot t^2$ and $\mathrm{qad}(\psi^\Phi_R) \leq \mathrm{qad}(\Phi) + 2$, for a constant c which depends only on σ, d.

Proof. Let $\Phi := \{\varphi_1, \dots, \varphi_\ell\}$. Consider some $i \in [\ell]$ and let $\tilde{\varphi}_i(x) := \varphi_i|_{\mathrm{reach}_d(x,z)}$.

Let $n \in [t]$. We define a formula $\psi^n_i(\bar{x})$, where $\bar{x} := (x_1, \dots, x_n)$, which states that $x_1, \dots, x_n$ lie in distinct connected components, each of which satisfies φ_i:

$$\psi^n_i(\bar{x}) \;:=\; \bigwedge_{j \in [n]} \tilde{\varphi}_i(x_j) \;\wedge\; \bigwedge_{j,k \in [n], j \neq k} \neg\, \mathrm{reach}_d(x_j, x_k).$$

Observe that $\mathrm{qad}(\psi^n_i) \leq \mathrm{qad}(\Phi)$ (in particular, since reach_d is an existential formula) and that $\|\psi^n_i\| \leq cn^2\|\Phi\| \leq ct^2\|\Phi\|$, for a constant c depending on σ, d only.

To obtain a formula which states that either the (pairwise disjoint) components of the $x_1, \dots, x_n$ are the only components which satisfy φ_i or the number of such components is at least t, we let

$$\psi^{n,t}_i(\bar{x}) \;:=\; \begin{cases} \forall y\, \neg\tilde{\varphi}_i(y) & \text{if } n = 0, \\ \psi^n_i(\bar{x}) \wedge \forall y\, (\tilde{\varphi}_i(y) \to \bigvee_{i \in [n]} \mathrm{reach}_d(y, x_i)) & \text{if } 0 < n < t \\ \psi^n_i(\bar{x}) & \text{if } n \geq t. \end{cases}$$

Note that $\mathrm{qad}(\psi_i^{n,t}) \leq \mathrm{qad}(\Phi) + 1$ and $\|\psi_i^{n,t}\| \leq c \cdot \|\psi_i^n\|$, for some constant c depending on σ, d only. (Note that $\|\psi_i^n\| \geq n$, so the disjunction over $i \in [n]$ is absorbed by that.) We obtain the desired sentence $\psi_{R,t}^\Phi$ by setting

$$\psi_{R,t}^\Phi := \bigvee_{(n_1,\dots,n_\ell)\in R} \exists \bar{x}_i \bigwedge_{i\in[\ell]} \psi_i^{n_i,t}(\bar{x}_i),$$

where $\bar{x}_i$ is a tuple of n_i variables. Note that

$$\begin{aligned} \|\psi_R^\Phi\| &\leq |R| \cdot |\Phi| \cdot \max_{i\in[\ell]} \|\psi_i^t\| \leq c \cdot |R| \cdot |\Phi| \cdot \|\Phi\| \cdot t^2, \\ \mathrm{qad}(\psi_R^\Phi) &\leq \max_{i\in[\ell]} \mathrm{qad}(\psi_i^{n_i,t}) + 1 \leq \mathrm{qad}(\Phi) + 2\,. \end{aligned}$$

□

Finally, we can prove our main theorem.

Proof of Theorem 4. By induction on the tree-depth d, we show that for each signature σ and each $\mathrm{FO}[\sigma^\leq]$-sentence φ with $\mathrm{qr}(\varphi) = q$, there is an $\mathrm{FO}[\sigma]$-sentence $\psi_{\varphi,d}$ with $\|\psi_{\varphi,d}\| \in d\text{-EXP}(q)$ and $\mathrm{qad}(\psi_{\varphi,d}) \leq 3d$ such that for each $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$, we have $\mathfrak{A} \models \psi_{\varphi,d}$ iff $\mathrm{tp}_q^\leq(\mathfrak{A}) \models \varphi$. Furthermore, we show that $|\mathcal{T}_{\sigma,q,d}| \in d\text{-EXP}(q)$ and $|\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}| \in (d-1)\text{-EXP}(q)$. To finish the proof, if φ is order-invariant, we let $\psi := \psi_{\varphi,d}$, and we obtain that $\mathfrak{A} \models_\leq \varphi$ iff $\mathfrak{A} \models \psi$.

Let $\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d} = \{\theta_1, \dots, \theta_\ell\}$. First, for each $i \in [\ell]$, we construct a sentence φ_i that defines θ_i on $\mathrm{Fin}^{\mathrm{conn}}_{\sigma,d}$. If $d = 1$, observe that any *connected* structure $\mathfrak{A}$ of type $\theta_i \in \mathcal{T}^{\mathrm{conn}}_{\sigma,q,1}$ consists of a single element. The atomic σ-type α of this element determines the q-type of the unique q-order on $\mathfrak{A}$. The FO-sentence $\varphi^{\mathrm{conn}}_{\tau,1} := \exists x\, \alpha(x)$ hence defines τ on $\mathrm{Fin}^{\mathrm{conn}}_{\sigma,1}$. We obviously have $\|\varphi^{\mathrm{conn}}_{\tau,1}\| \leq c \cdot |\sigma|$, for some absolute constant c, and $|\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}| \leq 2^{|\sigma|} \in (d-1)\text{-EXP}(q)$.

If $d > 1$, we construct an FO-sentence $\psi_{\theta,d-1}$ inductively for each q-type $\theta \in \mathcal{T}_{\tilde{\sigma},q,d-1}$. Let $\Psi := \{\psi_{\theta,d-1} \mid \theta \in \mathcal{T}_{\tilde{\sigma},q,d-1}\}$. By induction, we obtain $\|\Psi\| \in (d-1)\text{-EXP}(q)$, and $\mathrm{qad}(\Psi) \leq 3(d-1)$, and we have $|\mathcal{T}_{\tilde{\sigma},q,d-1}| \in (d-1)\text{-EXP}(q)$. We construct φ_i according to Lemma 10, i.e. we let $\varphi_i := \varphi^{\mathrm{conn}}_{\theta_i,d}$ for each $i \leq \ell$. Let $\Phi := \{\varphi_1, \dots, \varphi_\ell\}$. Then there is a constant c depending only on σ, d, such that

$$\begin{aligned} \|\Phi\| &\leq c \cdot \|\Psi\| \cdot |\mathcal{T}_{\tilde{\sigma},q,d-1}|^2 \in (d-1)\text{-EXP}(q) \quad \text{and} \\ \mathrm{qad}(\Phi) &\leq \mathrm{qad}(\Psi) + 2 \leq 3(d-1) + 2. \end{aligned}$$

Now consider a sentence $\varphi \in \mathrm{FO}[\sigma^\leq]$. Let $R := R_\varphi$ be given by Corollary 12. We apply Lemma 13 with $t := 2^q + 1$ to obtain a sentence $\psi_{\varphi,d} := \psi_R^\Phi$. To see that $\psi_{\varphi,d}$ is defined correctly, consider some $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$. Observe that for each $i \in [\ell]$ and each component $\mathfrak{K}$ of $\mathfrak{A}$, we have $\mathfrak{K} \models \varphi_i$ iff $\mathrm{tp}_q^\leq(\mathfrak{K}) = \tau_i$, and thus $\bar{n}_\Phi(\mathfrak{A}) = \bar{n}_{\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}}(\mathfrak{A})$. Then

$$\begin{aligned} \mathfrak{A} \models \psi_{\varphi,d} &\text{ iff } [\bar{n}_{\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}}(\mathfrak{A})]_{\wedge t} \in R && \text{(by Lemma 13 and previous observation)} \\ &\text{iff } \mathrm{tp}_q^\leq(\mathfrak{A}) \models \varphi. && \text{(by Corollary 12)} \end{aligned}$$

By Lemma 13, for some constant c depending only on σ, d, we have

$$\begin{aligned} \|\psi_{\varphi,d}\| &\leq c \cdot |\Phi| \cdot |R| \cdot t^2 \cdot \|\Phi\| \quad \text{and} \\ \mathrm{qad}(\psi_{\varphi,d}) &\leq \mathrm{qad}(\Phi) + 1 \leq 3d\,. \end{aligned}$$

Observe that $|\Phi| = \ell = |\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d}| \in (d-1)\text{-EXP}(q)$ by Corollary 8 and that $|R| \leq t^\ell \in d\text{-EXP}(q)$. Hence, $\|\psi_{\varphi,d}\| \in d\text{-EXP}(q)$. By Corollary 12, we also obtain $|\mathcal{T}_{\sigma,q,d}| \in d\text{-EXP}(q)$. □

4 Order-invariant monadic second-order logic

[5, Thm. 4.1] proved that classes of graphs definable by order-invariant MSO sentences are recognisable. Recognisable sets of graphs of *bounded tree-width* are conjectured in [4, Conjecture 1] to be definable in MSO with modulo-counting (CMSO), which would imply that <-inv-MSO is equivalent to CMSO on these graphs. Note that it is well-known and easy to see that, regardless of the considered class of structures, for each sentence of modulo-counting MSO there is an equivalent <-inv-MSO-sentence. Hence, the difficult part is the construction of an CMSO-sentence for a given <-inv-MSO-sentence.

While the equivalence of recognisability and definability in CMSO for graphs of bounded tree-width is still widely considered to be open (cf. [6, p. 574]), we show that in the further restricted case of structures of bounded tree-depth, <-inv-MSO collapses even to *first-order* logic with modulo counting (FO+MOD):

Theorem 14. *For every $d \in \mathbb{N}^+$ and every* <-inv-MSO*-sentence φ there is an* FO+MOD*-sentence ψ with $\mathrm{qad}(\psi) \leq 3d$ which is equivalent to φ on $\mathrm{Fin}_{\sigma,d}$.*

In contrast to the previous sections, we do not analyse the formula size, because it is known from [11] that (plain) MSO can define the length of orders non-elementarily more succinct than FO.

For the proof of Theorem 14, we proceed similarly to the last section. Again we need to understand <-inv-MSO's capabilities to count the number of components of a given q-type in q-ordered structures. However, this time we need to count not only up to some threshold, but also modulo some fixed divisor.

For $n \in \mathbb{N}$ and $p \in \mathbb{N}^+$, we let $[n]_{\mathrm{mod}\, p}$ denote the remainder of the division of n by p, and $\bar{n} := (n_1, \dots, n_\ell) \in \mathbb{N}^\ell$, we let $[\bar{n}]_{\mathrm{mod}\, p} := ([n_1]_{\mathrm{mod}\, p}, \dots, [n_\ell]_{\mathrm{mod}\, p})$. Similarly, we set $m \equiv_{\mathrm{mod}\, p} n$ if p divides $m - n$, and extend this notion to tuples $\bar{m}$ and $\bar{n}$ component-wise.

Below, we prove the following Lemma which shows that MSO inherits its component counting capabilities in q-ordered structures from its capabilities to distinguish orders of different lengths.

Lemma 15. *For each $q \in \mathbb{N}^+$, there is a $p \in \mathbb{N}^+$ such that for all q-ordered structures $(\mathfrak{A}, \preceq^{\mathfrak{A}})$ and $(\mathfrak{B}, \preceq^{\mathfrak{B}})$,*

$$\big(\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A}) \equiv_{\mathrm{mod}\, p} \bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B}) \text{ and } \bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A}) \equiv_{\wedge p} \bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B})\big) \implies (\mathfrak{A}, \preceq^{\mathfrak{A}}) \equiv_q^{\mathrm{MSO}} (\mathfrak{B}, \preceq^{\mathfrak{B}}).$$

In the following, we say that an ordered structure $(\mathfrak{A}, \preceq)$ is *component ordered*, if the order $\preceq$ is a sum of the orders on the components of $\mathfrak{A}$, i.e. for some enumeration $\mathfrak{K}_1, \dots, \mathfrak{K}_n$ of the components of $\mathfrak{A}$, we have $\preceq = \preceq|_{K_1} + \preceq|_{K_2} + \cdots + \preceq|_{K_n}$. Observe that q-ordered structures are also component ordered. It will be convenient to have some notation that allows us to treat component ordered structures similarly to words. Given two ordered structures $(\mathfrak{A}, \preceq^{\mathfrak{A}})$ and $(\mathfrak{B}, \preceq^{\mathfrak{B}})$, we let $(\mathfrak{A}, \preceq^{\mathfrak{A}}) \sqcup (\mathfrak{B}, \preceq^{\mathfrak{B}}) := (\mathfrak{A} \sqcup \mathfrak{B}, \preceq^{\mathfrak{A}} + \preceq^{\mathfrak{B}})$, where $\mathfrak{A} \sqcup \mathfrak{B}$ denotes the disjoint union of $\mathfrak{A}$ and $\mathfrak{B}$ and we consider $\preceq^{\mathfrak{A}}, \preceq^{\mathfrak{B}}$ as orders on the components of the disjoint union (via the inclusion mappings for $\mathfrak{A}, \mathfrak{B}$). Instead of $(\mathfrak{A}, \preceq^{\mathfrak{A}}) \sqcup (\mathfrak{B}, \preceq^{\mathfrak{B}})$, we also write $(\mathfrak{A}, \preceq^{\mathfrak{A}})(\mathfrak{B}, \preceq^{\mathfrak{B}})$. Like in the following definition, we often omit the order to make this notation less cluttered. For each component ordered structure $\mathfrak{A}$, we define its *i-th power* $\mathfrak{A}^i$ by $\mathfrak{A}^1 := \mathfrak{A}$ and $\mathfrak{A}^i := \mathfrak{A}^{i-1}\mathfrak{A}$ if $i > 1$.

The proof of Lemma 15 rests on the following Lemma.

Lemma 16 (Pumping Lemma)**.** *For each $q \in \mathbb{N}^+$, there is a number $p \in \mathbb{N}^+$ such that for all component ordered structures $\mathfrak{A}$ and all $r \in \mathbb{N}$, $i, j \in \mathbb{N}^+$,*

$$\mathfrak{A}^{r+ip} \equiv_q^{\mathrm{MSO}} \mathfrak{A}^{r+jp}.$$

Proof. Let $\mathcal{T}$ denote the (finite) set of q-types which are realised by component ordered σ-structures. We lift the disjoint union of ordered structures to $\mathcal{T}$ by defining $\mathrm{tp}_q(\mathfrak{A}) \sqcup \mathrm{tp}_q(\mathfrak{B}) := \mathrm{tp}_q(\mathfrak{A} \sqcup \mathfrak{B})$. The Composition Lemma (Lemma 7) shows that this operation is well-defined. It is also associative, so that $(\mathcal{T}, \sqcup)$ is a finite semigroup. Hence, there is a number p such that for each $\tau \in \mathcal{T}$, τ^p is idempotent (cf. e.g. [12]), i.e. $\tau^p = \tau^{ip}$ for each $i \in \mathbb{N}^+$. Then, for all $\mathfrak{A}, r, i, p$ as in the statement of the lemma, $\mathrm{tp}_q(\mathfrak{A})^{r+ip} = \mathrm{tp}_q(\mathfrak{A})^{r+jp}$, i.e. $\mathfrak{A}^{r+ip} \equiv_q^{\mathrm{MSO}} \mathfrak{B}^{r+jp}$. □

Proof of Lemma 15. Let $\mathcal{T}_{\sigma,q} = \{\tau_1, \dots, \tau_\ell\}$ with $\tau_i \prec_q \tau_j$ iff $i < j$. For each $i \in [\ell]$, fix a connected q-ordered structure $\mathfrak{K}_i$ whose type is $\mathrm{tp}_q(\mathfrak{K}_i) = \tau_i$. By repeated application of the Composition Lemma, we can assume without loss of generality that $\mathfrak{K} \cong \mathfrak{K}_i$ for each q-ordered component $\mathfrak{K}$ of $\mathfrak{A}$ or $\mathfrak{B}$ with $\mathrm{tp}_q(\mathfrak{K}) = \tau_i$. Let $n_i := n_{\tau_i}(\mathfrak{A})$ and let $m_i := n_{\tau_i}(\mathfrak{B})$ for each $i \in [\ell]$. By part 2 of Definition 1, we obtain

$$\mathfrak{A} \cong \mathfrak{K}_1^{n_1} \mathfrak{K}_2^{n_2} \cdots \mathfrak{K}_\ell^{n_\ell} \quad \text{and} \quad \mathfrak{B} \cong \mathfrak{K}_1^{m_1} \mathfrak{K}_2^{m_2} \cdots \mathfrak{K}_\ell^{m_\ell}.$$

For each $i \in [\ell]$, we have $n_{\tau_i}(\mathfrak{A}) \equiv_{\mathrm{mod}\, p} n_{\tau_i}(\mathfrak{B})$, i.e. there are $r_i \in [0, p-1]$ and $a_i, b_i \in \mathbb{N}$ such that $n_i = r_i + a_i p$ and $m_i = r_i + b_i p$. Furthermore, as $n_{\tau_i}(\mathfrak{A}) \equiv_{\wedge p} n_{\tau_i}(\mathfrak{B})$, we have $a_i > 0$ iff $b_i > 0$. By repeated application of the Pumping Lemma, we obtain

$$\mathfrak{K}_1^{n_1} \mathfrak{K}_2^{n_2} \cdots \mathfrak{K}_\ell^{n_\ell} \equiv_q^{\mathrm{MSO}} \mathfrak{K}_1^{r_1 + b_1 p} \mathfrak{K}_2^{r_2 + b_2 p} \cdots \mathfrak{K}_\ell^{r_\ell + b_\ell p} = \mathfrak{K}_1^{m_1} \mathfrak{K}_2^{m_2} \cdots \mathfrak{K}_\ell^{m_\ell}.$$

Hence, $\mathfrak{A} \equiv_q^{\mathrm{MSO}} \mathfrak{B}$. □

The next lemma is a modulo-counting analogue of Lemma 13.

Lemma 17. *For all $d, p \in \mathbb{N}^+$, each set of* FO+MOD$[\sigma]$*-sentences Φ, and each set $R \subseteq [0,p]^\ell \times [0, p-1]^\ell$, there is an* FO+MOD$[\sigma]$*-sentence χ_R^Φ such that for each $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$,*

$$\mathfrak{A} \models \chi_R^\Phi \quad \textit{iff} \quad ([\bar{n}_\Phi(\mathfrak{A})]_{\wedge p}, [\bar{n}_\Phi(\mathfrak{A})]_{\mathrm{mod}\, p}) \in R.$$

Furthermore, $\mathrm{qad}(\chi_R^\Phi) \leq \max\{\mathrm{qad}(\Phi) + 2, 2(d-1) + 1\}$.

In contrast to Lemma 13, the proof of Lemma 17 is not straightforward, because it is not obvious how modulo-counting quantifiers can be used to count the number of components satisfying a given FO+MOD-sentence. A remedy to this problem is provided by the following Lemma 3, which shows that the number of tree-depth roots of each component of a graph (and hence of a structure) can be bounded in terms of its tree-depth only.

Proof of Lemma 17. Let $\Phi = \{\varphi_1, \dots, \varphi_\ell\}$. For each $\bar{n} \in [0,p]^\ell$, let $\varphi_{\{\bar{n}\}}^\Phi$ be given by Lemma 13 for $t := p$, i.e. for each $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$, we have $\mathfrak{A} \models \varphi_{\{\bar{n}\}}^\Phi$ iff $[\bar{n}_\Phi(\mathfrak{A})]_{\wedge p} = \bar{n}$. Furthermore, $\mathrm{qad}(\varphi_{\{\bar{n}\}}^\Phi) \leq \mathrm{qad}(\Phi) + 2$. Below, for each $\bar{r} := (r_1, \dots, r_\ell) \in [0, p-1]^\ell$, $i \in [\ell]$, we construct a sentence $\chi_i^{\bar{r}}$ such that $\mathfrak{A} \models \chi_i^{\bar{r}}$ iff $n_{\varphi_i}(\mathfrak{A}) \equiv_{\mathrm{mod}\, p} r_i$. Furthermore, $\mathrm{qad}(\chi_i^{\bar{r}}) \leq \max\{\mathrm{qad}(\Phi) + 1, 2(d-1) + 2\}$. We can then define $\chi_R^\Phi := \bigvee_{(\bar{n},\bar{r}) \in R} \big(\varphi_{\{\bar{n}\}}^\Phi \wedge \bigwedge_{i \in [\ell]} \chi_i^{\bar{r}}\big)$. Obviously, $\mathrm{qad}(\chi_R^\Phi) \leq \max\{\mathrm{qad}(\Phi) + 2, 2(d-1) + 2\}$.

Consider some $\bar{r} := (r_1, \dots, r_\ell) \in [0, p-1]^\ell$, $i \in [\ell]$, and let $\varphi := \varphi_i$ and $r := r_i$. We define a formula $\varphi^{=k}(x)$, such that $\mathfrak{A} \models \varphi^{=k}(a)$, for $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$ and $a \in A$, iff a belongs to a component $\mathfrak{K}$

of $\mathfrak{A}$ such that $\mathfrak{K} \models \varphi$, $a \in \operatorname{roots}(\mathfrak{K})$, and $|\operatorname{roots}(\mathfrak{K})| = k$. Let $\tilde{\varphi}(x) := \varphi|_{\operatorname{reach}_d(x,z)}$, let $\tilde{\operatorname{roots}}_d(x) := \operatorname{roots}_d(x)|_{\operatorname{reach}_d(x,z)}(x)$, and let

$$\begin{aligned}\varphi^{=k}(x) := \tilde{\varphi}(x) \ \wedge \ & \tilde{\operatorname{roots}}_d(x) \\ \wedge \exists x_1 \dots \exists x_k \ \Big(& \bigwedge_{j\in[k]} \big(\tilde{\operatorname{roots}}_d(x_j) \wedge \operatorname{reach}_d(x_j, x)\big) \ \wedge \bigwedge_{j,j'\in[k],\, j\neq j'} x_j \neq x_{j'}\big) \\ \wedge \forall y \ \big(& \tilde{\operatorname{roots}}_d(y) \wedge \bigwedge_{j\in[k]} y \neq x_j\big) \to \bigwedge_{j\in[k]} \neg \operatorname{reach}_d(y, x)\Big).\end{aligned}$$

Observe that

$$\begin{aligned}\operatorname{qad}(\varphi^{=k}) \ &\leq \ \max\{\operatorname{qad}(\tilde{\varphi}), \operatorname{qad}(\tilde{\operatorname{roots}}_d) + 1, \operatorname{qad}(\operatorname{reach}_d) + 1\} \\ &\leq \ \max\{\operatorname{qad}(\varphi), 2(d-1)+1\}.\end{aligned}$$

Let the function f be defined as in Lemma 3 and let $b := f(d)$. Let $M \subseteq [0, p-1]^{b+1}$ be such that

$$(a_0, \dots, a_b) \in M \quad \text{iff} \quad \sum_{k\in[0,b]} k \cdot a_k \equiv_{\operatorname{mod} p} r.$$

Now we define our formula $\chi_i^{\bar{n}}$ as

$$\chi_i^{\bar{n}} \ := \bigvee_{(a_1,\dots,a_b)\in M} \ \bigwedge_{k\in[0,b]} \exists^{k\cdot a_k \ (\operatorname{mod} p)} x \ \varphi^{=k}(x)\,.$$

Obviously, $\operatorname{qad}(\chi_i^{\bar{n}}) \ \leq \ \max\{\operatorname{qad}(\varphi), 2(d-1)+1\} + 1$.

We show that the formula is defined correctly. Let $\mathfrak{A} \in \operatorname{Fin}_{\sigma,d}$. Recall that, according to Lemma 3, $|\operatorname{roots}(\mathfrak{K})| \leq b$ for each component $\mathfrak{K}$ of $\mathfrak{A}$. We partition the set H of components of $\mathfrak{A}$ into pairwise disjoint sets $H_0, \dots, H_b$ such that $\mathfrak{K} \in H_k$ iff $|\operatorname{roots}(\mathfrak{K})| = k$, for each $\mathfrak{K} \in H$. By definition of $\varphi^{=k}(x)$, the number of elements $a \in A$ such that $\mathfrak{A} \models \varphi^{=k}(a)$ equals $k \cdot |H_k|$. Hence, $\mathfrak{A} \models \chi_i^{\bar{r}}$ iff for some $(a_0, \dots, a_b) \in M$, we have $k \cdot |H_k| \equiv k \cdot a_k \ (\operatorname{mod} p)$ for each $k \in [0, b]$. This is true iff $n_\varphi(\mathfrak{A}) \equiv r \ (\operatorname{mod} p)$, since

$$n_\varphi(\mathfrak{A}) \ = \ \sum_{k\in[0,b]} k \cdot |H_k| \ \equiv_{\operatorname{mod} p} \ \sum_{k\in[0,b]} k \cdot a_k \equiv_{\operatorname{mod} p} r,$$

for $a_0, \dots, a_b \in [0, p-1]$ such that $|H_k| \equiv_{\operatorname{mod} p} a_k$ for each $k \in [0, b]$. □

With these preparations, the proof of Theorem 14 is very similar to the proof of Theorem 4.

Proof of Theorem 14. The proof proceeds by induction on the tree-depth d. We show that for each MSO$[\sigma, \leq]$-sentence φ with $\operatorname{qr}(\varphi) = q$, there is an FO+MOD$[\sigma]$-sentence $\psi_{\varphi,d}$ such that for each $\mathfrak{A} \in \operatorname{Fin}_{\sigma,d}$, we have $\mathfrak{A} \models \psi_{\varphi,d}$ iff $\operatorname{tp}_q^{\leq}(\mathfrak{A}) \models \varphi$. In particular, if φ is order-invariant, we let $\psi := \psi_{\varphi,d}$, and we obtain $\mathfrak{A} \models_{\leq} \varphi$ iff $\mathfrak{A} \models \psi := \psi_{\varphi,d}$.

Let $\mathcal{T}^{\operatorname{conn}}_{\sigma,q,d} = \{\theta_1, \dots, \theta_\ell\}$. We construct a sentence φ_i that defines θ_i on $\operatorname{Fin}^{\operatorname{conn}}_{\sigma,d}$, for each $i \in [\ell]$. If $d = 1$, the type of a connected structure of type θ_i is determined by the atomic σ-type α of its single element. We let $\varphi^{\operatorname{conn}}_{\tau,1} \ := \ \exists x\, \alpha(x)$. If $d > 1$, for each q-type $\theta \in \mathcal{T}_{\tilde{\sigma},q,d-1}$, we obtain an FO+MOD-sentence $\psi_{\theta,d-1}$ with $\operatorname{qad}(\psi_{\theta,d-1}) \leq 3(d-1)$.

We construct φ_i according to Lemma 10, i.e. we let $\varphi_i := \psi^{\mathrm{conn}}_{\theta_i,d}$ for each $i \leq \ell$. Let $\Phi := \{\varphi_1, \dots, \varphi_\ell\}$. Note that $\mathrm{qad}(\Phi) \leq 3(d-1)+2$.

Now consider a sentence $\varphi \in \mathrm{MSO}[\sigma, \leq]$. Let

$$R := \left\{ \left([\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B})]_{\wedge p}, [\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B})]_{\mathrm{mod}\, p} \right) \mid \mathfrak{B} \in \mathrm{Fin}_{\sigma,d}, \mathrm{tp}^{\leq}_q(\mathfrak{B}) \models \varphi \right\}$$

where p is given by the Pumping Lemma for q. We construct $\psi_{\varphi,d} := \psi^{\Phi}_R$ according to Lemma 17. In particular, $\mathrm{qad}(\psi_{\varphi,d}) \leq \mathrm{qad}(\Phi)+1 \leq 3d$. Consider some $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$. Observe that, for each component $\mathfrak{K}$ of $\mathfrak{A}$, we have $\mathfrak{K} \models \varphi_i$ iff $\mathrm{tp}^{\leq}_q(\mathfrak{K}) = \tau_i$. Hence, $([\bar{n}_\Phi(\mathfrak{A})]_{\wedge p}, [\bar{n}_\Phi(\mathfrak{A})]_{\mathrm{mod}\, p}) = ([\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A})]_{\wedge p}, [\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A})]_{\mathrm{mod}\, p})$. Thus

$$\mathfrak{A} \models \psi_{\varphi,d} \iff ([\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A})]_{\wedge p}, [\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A})]_{\mathrm{mod}\, p}) = ([\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B})]_{\wedge p}, [\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B})]_{\mathrm{mod}\, p})$$

for some structure $\mathfrak{B} \in \mathrm{Fin}_{\sigma,d}$ with $\mathrm{tp}^{\leq}_q(\mathfrak{B}) \models \varphi$. As a consequence of Lemma 15, this holds iff $\mathrm{tp}^{\leq}_q(\mathfrak{A}) \models \varphi$. □

5 Monadic second-order logic

In [8] it was proved that each MSO-definable class of finite graphs of bounded tree-depth is also FO-definable. Our approach towards the results of the previous section can be adapted to obtain another proof of this result which allows us to give an elementary upper bound on the size of the FO-sentence in terms of the quantifier-rank of the MSO-sentence. Throughout this section, we assume in all notation whose definition refers to a logic L that $\mathrm{L} = \mathrm{MSO}$. We let $\mathcal{T}_{\sigma,q,d} := \{\mathrm{tp}_q(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{Fin}_{\sigma,d}\}$ and let $\mathcal{T}^{\mathrm{conn}}_{\sigma,q,d} := \{\mathrm{tp}_q(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{Fin}^{\mathrm{conn}}_{\sigma,d}\}$.

Theorem 18. *Let $d \in \mathbb{N}^+$ and let σ be a signature. For each $\mathrm{MSO}[\sigma]$-sentence φ there is an $\mathrm{FO}[\sigma]$-sentence ψ with $\|\psi\| \in d\text{-}\mathrm{EXP}(\mathrm{qr}(\varphi))$ and $\mathrm{qad}(\psi) \leq 2d$ that is equivalent to φ on $\mathrm{Fin}_{\sigma,d}$.*

We also prove the following theorem in Section 5.2 below which shows that the upper bound of Theorem 18 is essentially optimal.

Theorem 19. *There is a signature σ such that for each $d \in \mathbb{N}^+$ there is an $\mathrm{MSO}[\sigma]$-sentence φ_d such that each $\mathrm{FO}[\sigma]$-sentence ψ_d that is $\mathrm{Fin}_{\sigma,d}$-equivalent to φ_d has size $\|\psi_d\| \geq \|\varphi_d\|\text{-}\mathrm{EXP}(0)$.*

5.1 From MSO to FO

Much of the proof of Theorem 18 follows the proof of Theorem 4, but we are spared of the complications that arose in connection with the ordering of structures. Overall, this makes the proof of Theorem 18 simpler. On the other hand, the proof of an analogue to Lemma 11 becomes somewhat more complicated.

Counting components In Lemma 11, we did not use the fact that we consider only structures of bounded tree-depth. Here naively ignoring the bounded tree-depth would cause the component counting threshold for MSO-sentences of quantifier-rank q to depend non-elementarily on q. We use the following lemma to avoid this.

Lemma 20. *Let $d, q \in \mathbb{N}^+$. There is a $t := t(d,q) \in d\text{-}\mathrm{EXP}(q)$ such that for all structures $\mathfrak{A}, \mathfrak{B} \in \mathrm{Fin}_{\sigma,d}$,*

$$\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A}) \equiv_{\wedge t} \bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B}) \quad \Longrightarrow \quad \mathfrak{A} \equiv^{\mathrm{MSO}}_q \mathfrak{B}.$$

Lemma 20 is an easy consequence of the following two lemmas.

Lemma 21. *Let $k \in \mathbb{N}^+$, $q \in \mathbb{N}$, and $t := 2^{kq}$. Let σ be a signature. For all structures $\mathfrak{A}, \mathfrak{B} \in \mathrm{Fin}_\sigma$ whose components each contain at most k elements,*

$$\bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A}) \equiv_{\wedge t} \bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{B}) \quad \Longrightarrow \quad \mathfrak{A} \equiv_q^{\mathrm{MSO}} \mathfrak{B}.$$

Lemma 22. *Let $d, q \in \mathbb{N}^+$ and let σ be a signature. Each structure $\mathfrak{A} \in \mathrm{Fin}_{\sigma,d}$ contains an induced substructure $\mathfrak{B}$ with $|B| \in d\text{-EXP}(q)$ and $\mathfrak{A} \equiv_q^{\mathrm{MSO}} \mathfrak{B}$. If $\mathfrak{A}$ is connected, there is such a structure $\mathfrak{B}$ with $|B| \in (d-1)\text{-EXP}(q)$.*

Before we prove Lemma 21 and Lemma 22, we show how to prove Lemma 20 with their help. The proof will also use the following variant of a standard composition lemma, which we take for granted (we use a variant for signatures with constants, where the constant symbols will be used in the proof of Lemma 21).

The definition of the disjoint union $\mathfrak{A} \sqcup \mathfrak{B}$ of structures $\mathfrak{A}$ and $\mathfrak{B}$ can be extended to signatures with constant symbols, if the constant symbols of $\mathfrak{A}$ and $\mathfrak{B}$ are disjoint.

Lemma 23 (Composition Lemma)**.** *Let $q \in \mathbb{N}$. Let σ_1, σ_2 be signatures which may contain constant symbols, where the constants in σ_1 and σ_2 are disjoint. If $\mathfrak{A}_1, \mathfrak{B}_1$ are σ_1-structures and $\mathfrak{A}_2, \mathfrak{B}_2$ are σ_2-structures such that $\mathfrak{A}_1 \equiv_q^{\mathrm{MSO}} \mathfrak{B}_1$ and $\mathfrak{A}_2 \equiv_q^{\mathrm{MSO}} \mathfrak{B}_2$, then*

$$\mathfrak{A}_1 \sqcup \mathfrak{A}_2 \equiv_q^{\mathrm{MSO}} \mathfrak{B}_1 \sqcup \mathfrak{B}_2.$$

Proof of Lemma 20. With the help of Lemma 22 and the Composition Lemma, we can assume without loss of generality that $\mathfrak{A}$ and $\mathfrak{B}$ contain only components of size at most $k \in (d-1)\text{-EXP}(q)$. Let $t := 2^{kq}$ as in Lemma 21. Then $t \in d\text{-EXP}(q)$ and hence the claim follows from Lemma 21. □

Proof of Lemma 21. For the proof, we consider signatures σ which contain constant symbols. In this case, the components of a σ-structure are not necessarily σ-structures, because they might not contain all constants. Let $T_{\sigma,q}$ denote the union of the sets of $(\mathrm{MSO}, \sigma', q)$-types over all signatures $\sigma' \subseteq \sigma$. For σ-structures $\mathfrak{A}, \mathfrak{B}$ and $q, t \in \mathbb{N}^+$, we write $\mathfrak{A} \approx_{q,t} \mathfrak{B}$ if $\bar{n}_{T_{\sigma,q}} \equiv_{\wedge t} \bar{n}_{T_{\sigma,q}}$.

By induction on q, we prove the stronger claim that for each signature σ *which may contain constant symbols* and all σ-structures $\mathfrak{A}$ and $\mathfrak{B}$ whose components each contain at most k elements,

$$\mathfrak{A} \approx_{q,t} \mathfrak{B} \quad \Longrightarrow \quad \mathfrak{A} \equiv_q^{\mathrm{MSO}} \mathfrak{B}.$$

Let $q = 0$. Since $\mathfrak{A} \approx_{q,1} \mathfrak{B}$, there exists a bijection f between the sets $M_\mathfrak{A}, M_\mathfrak{B}$ of components of $\mathfrak{A}, \mathfrak{B}$ which contain constants. Furthermore, this bijection preserves the 0-type of components, i.e. for each component $\mathfrak{K} \in M_\mathfrak{A}$ there exists a partial isomorphism $g_\mathfrak{K}$ whose domain and codomain are, respectively, the set of constants of $\mathfrak{K}$ and $f(\mathfrak{K})$. These partial isomorphisms can be extended to a partial isomorphism $g := \bigcup_{\mathfrak{K} \in M_\mathfrak{A}} g_\mathfrak{K}$ of $\mathfrak{A}$ and $\mathfrak{B}$ whose domain and codomain are, respectively, the set of constants of $\mathfrak{A}$ and $\mathfrak{B}$. Hence $\mathfrak{A} \equiv_0^{\mathrm{MSO}} \mathfrak{B}$.

For each $q \in \mathbb{N}$, let $t(q) := 2^{kq}$. Now let $q > 0$. We consider the case where $\mathfrak{A}$ and $\mathfrak{B}$ contain only components of a single q-type τ over some signature $\sigma' \subseteq \sigma$. The general case follows by an application of the Composition Lemma. By a further application of the Composition Lemma, we can assume that all components of $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic to a single structure $\mathfrak{K}$ of type τ. Now if $n_\tau(\mathfrak{A}) = n_\tau(\mathfrak{B})$, then $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic, so we are done. Assume that $n_\tau(\mathfrak{A}), n_\tau(\mathfrak{B}) > t(q)$. We show that Duplicator wins the q-round EF-game on $\mathfrak{A}$ and $\mathfrak{B}$.

Consider the first round of the game. Suppose that Spoiler plays a point move, i.e. he chooses an element, say, $a \in A$. Duplicator chooses an element b corresponding to a in a copy of $\mathfrak{K}$ in $\mathfrak{B}$. This introduces exactly one component of a new isomorphism-type τ' in each of $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$. The remaining components of $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ all remain their isomorphism-type and there are more than $t(q) - 1 \geq t(q-1)$ such components. Hence $(\mathfrak{A}, a) \approx_{q-1,t(q-1)} (\mathfrak{B}, b)$. By induction, $(\mathfrak{A}, a) \equiv^{\text{MSO}}_{q-1} (\mathfrak{B}, b)$. So Duplicator wins, if she replies by b.

Suppose now that Duplicator plays a set move, say, $M \subseteq A$. Since $\mathfrak{K}$ contains at most k elements, the components of the structure $(\mathfrak{A}, M)$ belong to at most 2^k different isomorphism-types. Thus the number of q-types cannot be greater either. For each q-type θ occurring in $(\mathfrak{A}, M)$, let C_θ denote the set of components of $\mathfrak{A}$ whose q-type is θ. Duplicator chooses a set C'_θ of components of $\mathfrak{B}$ and a set of elements $M'_\theta \subseteq \bigcup_{\mathfrak{C} \in C'_\theta} C$ such that $\min\{|C_\theta|, t(q-1)\} = \min\{|C'_\theta|, t(q-1)\}$, and $\operatorname{tp}_q(\mathfrak{C}, M'_\theta \cap C) = \theta$ for each $\mathfrak{C} \in C'_\theta$. Since there are $t(q) > 2^k \cdot t(q-1)$ copies of $\mathfrak{K}$ in $\mathfrak{B}$, this is possible. Let $M' := \bigcup_\theta M'_\theta$. We have $(\mathfrak{A}, M) \approx_{q-1,t(q-1)} (\mathfrak{B}, M')$. So, by induction, $(\mathfrak{A}, M) \equiv^{\text{MSO}}_{q-1} (\mathfrak{B}, M')$. Replying by M', Duplicator wins. □

Lemma 22 is an adaptation of [16, Thm. 6.7] from FO to MSO. Its proof uses the previous lemma and the following analogue to Lemma 1, which can be proved like Lemma 1.

Lemma 24. *Let $q \in \mathbb{N}^+$. Let $\mathfrak{A}, \mathfrak{B} \in \operatorname{Fin}_\sigma$ be connected structures with $\operatorname{td}(\mathfrak{A}), \operatorname{td}(\mathfrak{B}) > 1$ and let $r_\mathfrak{A} \in \operatorname{roots}(\mathfrak{A})$,$r_\mathfrak{B} \in \operatorname{roots}(\mathfrak{B})$ with $\alpha(\mathfrak{A}, r_\mathfrak{A}) = \alpha(\mathfrak{B}, r_\mathfrak{B})$. Then*

$$\mathfrak{A}^{[r_\mathfrak{A}]} \equiv^{\text{MSO}}_q \mathfrak{B}^{[r_\mathfrak{B}]} \implies \mathfrak{A} \equiv^{\text{MSO}}_q \mathfrak{B}.$$

Proof of Lemma 22. The proof is by induction on the tree-depth d. First, we consider the claim about connected structures. If $d = 1$, then each connected structure with $\operatorname{td}(\mathfrak{A}) = 1$ has size $1 \in 0\text{-EXP}(q)$, i.e. we can set $\mathfrak{B} := \mathfrak{A}$. Suppose now that $d > 1$. Choose a tree-depth root $r \in \operatorname{roots}(\mathfrak{A})$. By induction, since $\operatorname{td}(\mathfrak{A}^{[r]}) \leq d - 1$, we obtain an induced substructure $\mathfrak{B}'$ of $\mathfrak{A}^{[r]}$ such that $|B'| \in (d-1)\text{-EXP}(q)$ and $\mathfrak{B}' \equiv^{\text{MSO}}_q \mathfrak{A}^{[r]}$. Let $\mathfrak{B}$ be the substructure of $\mathfrak{A}$ induced by $B' \cup \{r\}$, i.e. $\mathfrak{B}^{[r]} = \mathfrak{B}'$. Since $\mathfrak{A}^{[r]} \equiv^{\text{MSO}}_q \mathfrak{B}^{[r]}$, we obtain that $\mathfrak{A} \equiv^{\text{MSO}}_q \mathfrak{B}$ in the same way as in Lemma 1. Observe that $|B| \in (d-1)\text{-EXP}(q)$.

Consider the case that $\mathfrak{A}$ is not connected. By the construction above, we can replace each component $\mathfrak{K}$ of $\mathfrak{A}$ by an induced substructure of $\mathfrak{K}$ on $(d-1)\text{-EXP}(q)$ vertices that has the same q-type as $\mathfrak{K}$. By the Composition Lemma, this preserves the q-type of $\mathfrak{A}$. Let $k \in (d-1)\text{-EXP}(q)$ denote the maximum number of vertices in a component of $\mathfrak{A}$ after this replacement. By Lemma 21, we know that $\mathfrak{B} \equiv^{\text{MSO}}_q \mathfrak{A}$ for each induced substructure $\mathfrak{B}$ of $\mathfrak{A}$ such that $n_\tau(\mathfrak{B}) \equiv_{\wedge t} n_\tau(\mathfrak{A})$ for each q-type τ, where $t := 2^{kq}$. Since there are at most 2^k non-isomorphic components in $\mathfrak{A}$ and we have to keep at most t copies of each such component, there is such a structure $\mathfrak{B}$ with $|B| \in d\text{-EXP}(q)$. □

Finishing the proof With the preparations above, the proof of Theorem 18 is now very similar to the proof of Theorem 4.

Proof of Theorem 18. The proof proceeds by induction on the tree-depth d, where we also show that $|\mathcal{T}_{\sigma,q,d}| \in d\text{-EXP}(q)$ and $|\mathcal{T}^{\text{conn}}_{\sigma,q,d}| \in (d-1)\text{-EXP}(q)$.

Defining types of connected structures As a first step, we prove that each q-type $\tau \in \mathcal{T}^{\text{conn}}_{\sigma,q,d}$ is $\text{Fin}^{\text{conn}}_{\sigma,d}$-equivalent to an $\text{FO}[\sigma]$-sentence $\varphi^{\text{conn}}_{\tau,d}$ such that $\|\varphi^{\text{conn}}_{\tau,d}\| \in (d-1)\text{-EXP}(q)$ and $\text{qad}(\varphi^{\text{conn}}_{\tau,d}) \leq 3(d-1)+1$. For $d = 1$, each structure $\mathfrak{A} \in \text{Fin}^{\text{conn}}_{\sigma,d}$ of type τ consists of a single element of some atomic σ-type α. The FO-sentence $\varphi^{\text{conn}}_{\tau,1} := \exists x\, \alpha(x)$ then defines τ. Hence $\|\varphi^{\text{conn}}_{\tau,1}\|$ does not depend on q, $\text{qad}(\varphi^{\text{conn}}_{\tau,1}) = 0$, and $|\mathcal{T}^{\text{conn}}_{\sigma,q,d}| \leq 0\text{-EXP}(q)$.

Now suppose that $d > 1$ and let $\tau \in \mathcal{T}^{\text{conn}}_{\sigma,q,d}$. Let $R \subseteq \mathcal{T}_{\tilde{\sigma},q,d-1} \times 2^{\sigma}$ be a set that contains (θ, α) iff there is a structure $\mathfrak{B} \in \text{Fin}^{\text{conn}}_{\sigma,d}$ with $\text{tp}_q(\mathfrak{B}) = \tau$ which contains a tree-depth root $r \in \text{roots}(\mathfrak{B})$ such that $\alpha(\mathfrak{B}, r) = \alpha$ and $\text{tp}_q(\mathfrak{B}^{[r]}) = \theta$. Observe that, as a consequence of Lemma 24, for each $\mathfrak{A} \in \text{Fin}^{\text{conn}}_{\sigma,d}$, we have $\text{tp}_q(\mathfrak{A}) = \tau$ iff $(\text{tp}_q(\mathfrak{A}^{[r]}), \alpha(\mathfrak{A}, r)) \in R$ for some $r \in \text{roots}(\mathfrak{A})$. Now consider a q-type $\theta \in \mathcal{T}^{\text{conn}}_{\tilde{\sigma},q,d-1}$ and let $\varphi_{\theta,d-1}$ be the $\text{FO}[\tilde{\sigma}]$-sentence, given by induction, which is equivalent to θ on $\text{Fin}^{\text{conn}}_{\tilde{\sigma},d-1}$. As a consequence of Lemma 2, we obtain that for all structures $\mathfrak{A} \in \text{Fin}^{\text{conn}}_{\sigma,d}$ with $\text{td}(\mathfrak{A}) > 1$ and all tree-depth roots $r \in \text{roots}(\mathfrak{A})$, we have $\mathfrak{A} \models \mathcal{I}(\varphi_{\theta,d-1})(r)$ iff $\text{tp}_q(\mathfrak{A}^{[r]}) = \theta$.

Altogether, we obtain that the following $\text{FO}[\sigma]$-sentence is equivalent to τ on $\text{Fin}^{\text{conn}}_{\sigma,d}$:

$$\varphi^{\text{conn}}_{\tau,d} \;:=\; (\text{td} \leq 1 \wedge \varphi_{\tau,d-1}) \;\vee \bigvee_{(\theta,\alpha)\in R} \exists x \,\big(\text{roots}_d(x) \;\wedge\; \alpha(x) \;\wedge\; \mathcal{I}(\varphi_{\theta,d-1})(x)\big)\,.$$

Recall that, by induction, $\|\mathcal{I}(\varphi_{\theta,d-1})\| \in (d-1)\text{-EXP}(q)$ and $|\mathcal{T}_{\tilde{\sigma},q,d-1}| \in (d-1)\text{-EXP}(q)$. Hence, $|R| \in (d-1)\text{-EXP}(q)$. Altogether, we obtain that $\|\varphi^{\text{conn}}_{\tau,d}\| \in (d-1)\text{-EXP}(q)$. Using Lemma 24, we conclude that $|\mathcal{T}^{\text{conn}}_{\sigma,q,d}| \leq 2^{\sigma} \cdot |\mathcal{T}_{\tilde{\sigma},q,d-1}| \in (d-1)\text{-EXP}(q)$. By induction, $\text{qad}(\mathcal{I}(\varphi_{\theta,d-1})) \leq 3(d-1)$. Hence, $\text{qad}(\varphi^{\text{conn}}_{\tau,d}) \leq 3(d-1)+1$.

Structures with multiple components Consider an $\text{MSO}[\sigma]$-sentence φ. Let $\mathcal{T}^{\text{conn}}_{\sigma,q,d} := \{\tau_1, \ldots, \tau_\ell\}$, where $\ell := |\mathcal{T}^{\text{conn}}_{\sigma,q,d}|$. Let $t := t(d,q) \in d\text{-EXP}(q)$ be given by Lemma 20. Let Φ be the set that contains the formulae $\varphi_i := \varphi^{\text{conn}}_{d,\tau_i}$ for each $i \in [\ell]$. Hence, $\bar{n}_\Phi(\mathfrak{A}) \equiv_{\wedge t} \bar{n}_{\mathcal{T}^{\text{conn}}_{\sigma,q,d}}(\mathfrak{A}) = \bar{n}_{\mathcal{T}_{\sigma,q}}(\mathfrak{A})$ for each $\mathfrak{A} \in \text{Fin}_{\sigma,d}$. Let $R \subseteq [0,t]^\ell$ be a set such that $\bar{n} \in R$ iff there exists a model $\mathfrak{A} \in \text{Fin}_{\sigma,d}$ of φ with $[\bar{n}_\Phi(\mathfrak{A})]_{\wedge t} = \bar{n}$. Using Lemma 20, we obtain that $\mathfrak{A} \models \varphi$ iff $[\bar{n}_\Phi(\mathfrak{A})]_{\wedge t} \in R$, for each $\mathfrak{A} \in \text{Fin}_{\sigma,d}$. Hence, the $\text{FO}[\sigma]$-sentence $\psi := \psi^{\Phi}_{R}$ of Lemma 13 is equivalent to ψ on $\text{Fin}_{\sigma,d}$.

Regarding the size of ψ, note that Lemma 20 implies that $|R| \;\leq\; |\mathcal{T}^{\text{conn}}_{\sigma,q,d}| \;\leq\; [0,t]^\ell$. Since

$$\begin{aligned} t^\ell \;\in\; (d\text{-EXP}(q))^{(d-1)\text{-EXP}(q)} \;&= (2^{(d-1)\text{-EXP}(q)})^{(d-1)\text{-EXP}(q)} \\ &= 2^{(d-1)\text{-EXP}(q)\cdot(d-1)\text{-EXP}(q)} \\ &\subseteq 2^{(d-1)\text{-EXP}(q)} \;=\; d\text{-EXP}(q) \end{aligned}$$

we obtain that, by the construction of ψ according to Lemma 13,

$$\begin{aligned} \|\psi\| \;&\leq c \cdot |\Phi| \cdot \|\Phi\| \cdot |R| \cdot t^2. \\ &\in (d-1)\text{-EXP}(q) \cdot d\text{-EXP}(q) \cdot d\text{-EXP}(q)^2 \cdot (d-1)\text{-EXP}(q) \\ &\subseteq d\text{-EXP}(q), \end{aligned}$$

and $\text{qad}(\psi) \leq \text{qad}(\Phi) + 2 \leq 3d$. □

5.2 A lower bound

The proof of Theorem 19 uses an encoding of large natural numbers n by shallow trees $\mathrm{enc}(n)$ from [9, chapter 10.3]. Here, by *trees*, we mean directed trees which are rooted, i.e. trees which contain a root vertex from which all edges point away. The encoding is defined inductively as follows:

- $\mathrm{enc}(0)$ is the one-node tree.
- For $n \geq 1$, the tree $\mathrm{enc}(n)$ is obtained by creating a new root and attaching to it all trees $\mathrm{enc}(i)$ such that the i-th bit in the binary representation of n is 1.

Note that a tree encodes a number with respect to this encoding iff there are no two distinct isomorphic subtrees whose roots are children of the same vertex. But we would like to assign a natural number to each tree. To this end, we reduce each tree $\mathfrak{T}$ in a bottom-up way to a tree $\mathrm{num}(\mathfrak{T})$ that encodes a number:

- $\mathrm{num}(\mathfrak{T}) := \mathfrak{T}$ if $\mathrm{height}(\mathfrak{T}) = 1$, i.e. $\mathfrak{T} \cong \mathrm{enc}(0)$.
- If $\mathrm{height}(\mathfrak{T}) > 1$, select one tree $\mathfrak{T}_1, \ldots, \mathfrak{T}_k$ of each isomorphism type that occurs among the immediate subtrees of the root of $\mathfrak{T}$. Define $\mathrm{num}(\mathfrak{T})$ to be a tree whose root has children whose rooted subtrees are $\mathrm{num}(\mathfrak{T}_1), \ldots, \mathrm{num}(\mathfrak{T}_k)$.

Throughout the following section, we let $\sigma := \{E, R, B\}$, where E is a binary and R, B are unary relation symbols. We consider a tree as a $\{E\}$-structure $\mathfrak{T}$ where $E^{\mathfrak{T}}$ is the edge relation of the tree. A *coloured tree* is a finite σ-structure $(\mathfrak{T}, R^{\mathfrak{T}}, B^{\mathfrak{T}})$, where $\mathfrak{T}$ is a tree and $R^{\mathfrak{T}}, B^{\mathfrak{T}}$ (the *red* and the *blue* vertices of $\mathfrak{T}$) form a partition of the vertex set of the tree. Structures whose components are (coloured) trees are called (coloured) *forests*. The *height* $\mathrm{height}(\mathfrak{T})$ of a (coloured) tree $\mathfrak{T}$ is the maximum number of vertices on a path from the root of $\mathfrak{T}$ to a leave of $\mathfrak{T}$. The *height* $\mathrm{height}(\mathfrak{F})$ of a (coloured) forest $\mathfrak{F}$ is the maximum height of its components.

From the proof of [9, Lemma 10.21][2], we obtain the following lemma.

Lemma 25. *For each $d \in \mathbb{N}^+$, there is an $\mathrm{FO}[E]$-formula $eq_d(x, y)$ of size $\|eq_d\| \in O(d)$ such that for all forests $\mathfrak{F}$ with $\mathrm{height}(\mathfrak{F}) \leq d$ and all trees $\mathfrak{T}_1, \mathfrak{T}_2$ of $\mathfrak{F}$ with roots u_1, u_2, respectively, we have:*

$$\mathfrak{F} \models eq_d(u_1, u_2) \iff \mathrm{num}(\mathfrak{T}_1) = \mathrm{num}(\mathfrak{T}_2).$$

Note that $\mathrm{height}(\mathrm{enc}(n)) \leq d$ provided that $n < \mathrm{tower}(d)$, where $\mathrm{tower}(d) := d\text{-}\mathrm{EXP}(0)$. For each $d \geq 1$, let $\mathfrak{F}_d$ denote a coloured forest that contains exactly the trees $\mathrm{enc}(0), \ldots, \mathrm{enc}(\mathrm{tower}(d) - 1)$ whose vertices all are coloured red, let $\mathfrak{T}_d$ denote a coloured tree with $\mathrm{height}(\mathfrak{T}_d) \leq d$ that contains each of the trees $\mathrm{enc}(0), \ldots, \mathrm{enc}(\mathrm{tower}(d-1) - 1)$? as subtrees (e.g. a full $\mathrm{tower}(d-1)$-ary tree) and where all vertices are blue, and let $\mathfrak{F}_d^n$ denote the disjoint union of $\mathfrak{F}_d$ and n disjoint copies of $\mathfrak{T}_d$, for each $n \geq 0$.

Lemma 26. *For each $d \in \mathbb{N}^+$, there exists an $\mathrm{MSO}[\sigma]$-sentence φ_d of size $O(d)$ such that $\mathfrak{F}_d^n \models \varphi_d$ iff $n \geq \mathrm{tower}(d)$.*

[2] [9, Lemma 10.21] makes the assumption that $\mathfrak{T}_1, \mathfrak{T}_2$ are encodings of numbers n, m to conclude that $\mathfrak{F} \models \mathrm{eq}_d(u_1, u_2) \iff n = m$, i.e. $\mathfrak{T}_1 \cong \mathfrak{T}_2$. If we drop this assumption, we obtain our variant of the lemma using exactly the same formula.

Proof. Let $d \in \mathbb{N}^+$ and let $\mathrm{eq}_d(x, y, M)$ be the relativisation of the FO$[E]$-formula of Lemma 25 to a set variable M. Let $\mathrm{conn}(M)$ be an MSO$[E]$-formula which states in a forest $\mathfrak{F}$ that for each tree $\mathfrak{T}$ of $\mathfrak{F}$, the structure induced by M in $\mathfrak{T}$ is connected, i.e. a tree. Let $\mathrm{root}(x, M)$ state that x is a root in the subforest induced by M. We can assume that the size of $\mathrm{conn}(M)$ and $\mathrm{root}(x, M)$ is independent of d. Now let φ_d be the following sentence:

$$\exists M \Big(\mathrm{conn}(M) \ \wedge \ \forall x \ \big(R(x) \wedge \mathrm{root}(x, M)\big) \rightarrow \\ \exists y \ \big(\mathrm{root}(y, M) \ \wedge \ B(y) \ \wedge \ \mathrm{eq}_h(x, y, M)\big)\Big).$$

First we argue that $n \geq \mathrm{tower}(h)$ implies $\mathfrak{F}_d^n \models \varphi_d$. By definition, the red trees contained in $\mathfrak{F}_d^n$ are $\mathrm{enc}(0), \ldots, \mathrm{enc}(\mathrm{tower}(d) - 1)$. Since $n \geq \mathrm{tower}(h)$, we can choose $\mathrm{tower}(h)$ pairwise distinct copies $\mathfrak{H}_0, \ldots, \mathfrak{H}_{\mathrm{tower}(h)-1}$ of $\mathfrak{T}_d$ in $\mathfrak{F}_d^n$. Since all trees $\mathrm{enc}(0), \ldots, \mathrm{enc}(\mathrm{tower}(d)-1)$ occur as subtrees of $\mathfrak{T}_d$, for each $i \in [0, \mathrm{tower}(d) - 1]$ there is a set $M_i \subseteq H_i$ with $(\mathfrak{H}_n[M_n])|_E \cong \mathrm{enc}(i)$. The set $M := M_1 \cup \cdots \cup M_n$ witnesses that $\mathfrak{F}_d^n \models \varphi_d$.

Now we show that $\mathfrak{F}_d^n \models \varphi_d$ implies $n \geq \mathrm{tower}(h)$. Let $M \subseteq F_d^n$ witness that $\mathfrak{F}_d^n \models \varphi_d$. The forest $\mathfrak{F}_d^n$ contains trees $\mathrm{enc}(0), \ldots, \mathrm{enc}(\mathrm{tower}(d) - 1)$ whose vertices are all red. Hence, and according to the choice of M and the choice of $\mathrm{eq}_h(x, y, M)$, for each $i \in [0, \mathrm{tower}(d) - 1]$ there is a blue copy $\mathfrak{T}$ of $\mathfrak{T}_d$ in $\mathfrak{F}_d^n$ such that $\mathrm{num}(\mathfrak{T}[M]) = \mathrm{num}(\mathrm{enc}(i)) = i$. Hence $\mathfrak{F}_d^n$ must contain at least $\mathrm{tower}(h)$ copies of $\mathfrak{T}_d$, because M induces at most one tree in each copy of T_h. □

Using Lemma 26, we can easily finish the proof of Theorem 19.

Proof of Theorem 19. FO-sentences of quantifier-rank q cannot distinguish $\mathfrak{F}_d^k$ from $\mathfrak{F}_d^{k+1}$ for each $k \geq q$. Hence an FO-sentence ψ_d that is equivalent to the MSO-sentence φ_d of Lemma 26 must have quantifier-rank $\mathrm{qr}(\psi_d) \geq \mathrm{tower}(d)$ and in particular $\|\psi_d\| \geq \mathrm{tower}(d)$. □

6 Defining Bounded-Depth Tree-Decompositions in FO

For every finite relational signature σ and every $k \in \mathbb{N}$ there is a set $\Sigma(\sigma, k)$ of labels such that information about a σ-structure $\mathfrak{A}$ of tree-width at most k may be encoded into a $\Sigma(\sigma, k)$-labelled tree $T_{\mathfrak{A}}$. This encoding may be chosen so that the original structure $\mathfrak{A}$ can be interpreted in $T_{\mathfrak{A}}$ by an MSO-interpretation. One such encoding is presented in details in [9, Section 11.4].[3]

The question of whether there is an interpretation in the converse direction, i.e. whether some tree $T_{\mathfrak{A}}$ representing a width-k tree-decomposition of $\mathfrak{A}$ can be MSO-interpreted in $\mathfrak{A}$, is still open. In particular, interpretability of such a decomposition would imply that recognisability equals CMSO-definability for graphs of bounded tree-width.

In this section we show that for graphs of bounded tree-*depth*, there is even an FO-interpretation of a bounded-depth tree-decomposition. Since the interpretation we give here is not parameterised we obtain a canonical tree-decomposition, though not one of optimal depth or width. The FO-interpretation is given by formulae $\epsilon_d(x, y)$ and $\alpha_d(x, y)$ for every $d \geq 1$ such that if $\mathfrak{A}$ is a σ-structure of tree-depth at most d then

- ϵ_d defines an equivalence relation $\sim_{\mathfrak{A}} := \{(u, v) \mid \mathfrak{A} \models \epsilon_d[u, v]\}$ on A,
- the equivalence classes of $\sim_{\mathfrak{A}}$ have size bounded by a function of d,

[3] That $\mathfrak{A}$ can be MSO-interpreted in $T_{\mathfrak{A}}$ is not proved there but easy to see.

- the relation defined by α_d is invariant under $\sim_{\mathfrak{A}}$, i.e. if $u \sim_{\mathfrak{A}} u'$ and $v \sim_{\mathfrak{A}} v'$, then
$$\mathfrak{A} \models \alpha_d(u,v) \iff \mathfrak{A} \models \alpha_d(u',v'), \text{ and}$$
- α_d defines a rooted tree structure on the quotient structure $\mathfrak{A}/\sim_{\mathfrak{A}}$, in which $[u]_{\sim_{\mathfrak{A}}}$ is an ancestor of $[v]_{\sim_{\mathfrak{A}}}$ or vice versa whenever $u, v \in A$ are adjacent in the Gaifman graph of $\mathfrak{A}$.

This can be turned into a bounded-depth tree-decomposition in the usual sense by taking the tree structure on $\mathfrak{A}/\sim_{\mathfrak{A}}$ as the tree and setting $\{v \mid [v]_{\sim_{\mathfrak{A}}} \text{ is an ancestor of } [u]_{\sim_{\mathfrak{A}}}\}$ as the bag of the node $[u]_{\sim_{\mathfrak{A}}}$.

The key insight we use is Lemma 3 which says that for any fixed d there are at most $f(d)$ many candidates which may be placed at the root of a tree-decomposition of $\mathfrak{A}$ of minimum height. We have already seen at the end of Section 2 that there is an FO-formula $\text{roots}_d(x)$ such that $\mathfrak{A} \models \text{roots}_d[r]$ iff r is such a candidate. We recursively build a tree-decomposition $\mathcal{T}_{\mathfrak{A}}$ of $\mathfrak{A}$ of height at most d by placing, in each step, *all* candidate roots into the root-bag of our tree-decomposition and then recursing on the components of the remaining graph. Note that even if $\text{td}(\mathfrak{A}) = d$, not all components of $\mathfrak{A} \setminus R$, where R is the set of at most $f(d)$ root nodes, necessarily have tree-depth $d-1$, so we must be a bit careful which elements we place into the root of the next level.

We fix a tree-depth d and recursively define FO-formulae φ_i for $i = 0, \dots, d$ with the intended meaning that, in a structure $\mathfrak{A}$ of tree-depth d with $a \in A$, $\mathfrak{A} \models \varphi_i[a]$ iff a is on the i-th level of the tree-decomposition, which we denote by L_i:

$$\begin{aligned}\varphi_0(x) &:= \bot \\ \varphi_i(x) &:= \bigvee_{j=1}^{d-i} \left(\text{td}_{=j+1} \mid_{\neg\varphi_{<i}} \wedge \text{td}_{=j} \mid_{\neg(\varphi_{<i} \vee z \doteq x)} \right)\end{aligned}$$

Here, x is the free variable of φ_i and z is the free variable of the formulae used in the restrictions. With the abbreviations

$$\varphi_{<i}(x) := \bigvee_{j<i} \varphi_j(x) \quad \text{and} \quad \varphi_{\leq i}(x) := \bigvee_{j \leq i} \varphi_j(x)$$

we define

$$\begin{aligned}\psi_0(x,y) &:= \top \\ \psi_{i+1}(x,y) &:= \text{reach}_{d-i+1} \mid_{\neg\varphi_{\leq i}},\end{aligned}$$

i.e. $\psi_i(u,v)$ holds iff u and v are in the same connected component of $\mathfrak{A} \setminus \bigcup_{j \leq i} L_j$. We can now define an equivalence relation on $\mathfrak{A}$ as follows:

$$\epsilon_d(x,y) := \bigvee_{1 \leq i \leq d} (\varphi_i(x) \wedge \varphi_i(y) \wedge \psi_i(x,y)),$$

i.e. two elements are equivalent iff they appear on the same level of our tree-decomposition and are in the same connected component of $\mathfrak{A}$ after removing the levels above x and y. This is equivalent to saying that x and y appear in the same node of our tree-decomposition.

Let $\gamma(x,y)$ be a formula which expresses that to elements are adjacent in the Gaifman graph of a structure. Finally, We define tree edges (directed towards the root) by

$$\alpha_d(x,y) := \bigvee_{1 \leq i < d} (\varphi_i(x) \wedge \varphi_{i+1}(y) \wedge \exists u \exists v \, (\gamma(u,v) \wedge \epsilon(x,u) \wedge \psi_{i+1}(y,v))).$$

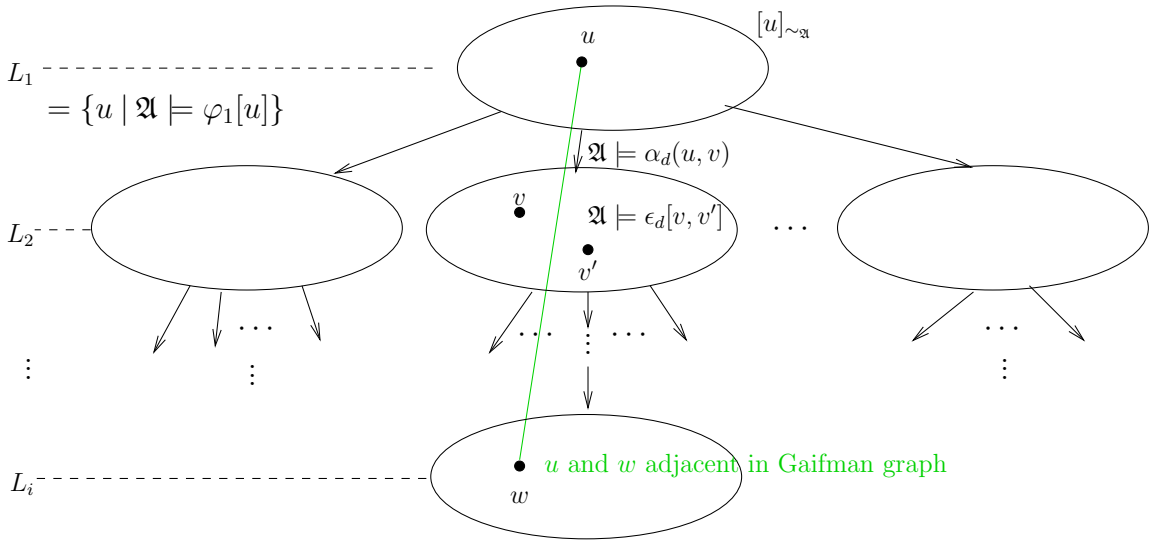

Figure 2: The canonical tree-decomposition defined in FO.

7 Conclusion

We have investigated the expressive power and the relative succinctness of different classes of logical formulae on structures of bounded tree-depth d. In particular, we have shown that, if a class $\mathcal{C}$ of such structures is MSO-definable or order-invariantly FO-definable, then it is also FO-definable. For MSO-definable classes, this was already known. But, in both cases, our approach also shows that the size of the FO-sentence which defines $\mathcal{C}$ is at most d-fold exponential in the quantifier-rank of a given order-invariant FO- or MSO-sentence which defines $\mathcal{C}$. For MSO-formulae, we have proved that this upper bound on the size of the FO-sentence is essentially optimal. It would be interesting to know if there is a corresponding lower bound for the result about order-invariantly FO-definable classes.

One motivation to consider bounded tree-depth graphs was the role of these graphs in the theory of sparse graphs which has been outlined in the book [16]. This link has been exploited in several results about the algorithmic behaviour of logics on sparse structures. Can our results on order-invariant FO-sentences on bounded tree-depth structures be used to obtain results about such sentences on more general classes of sparse structures?

An interesting extension of order-invariance is *addition-invariance* where sentences are not only allowed to use some linear order but also the graph of the addition operation that is induced by the embedding of a structure into the natural numbers that comes with the linear order. The paper [17] obtained a characterisation of the classes of structures which are addition-invariantly FO-definable over unary signatures, i.e. on structures of tree-depth 1. Each such class of structures is definable in FO_{card}, i.e. the extension of FO with nullary predicates C_m, for all positive integers m, which state that the cardinality of a structure is divisible by m. Our proofs hinge on the composition method and there is no obvious way how these methods could be extended to addition-invariant formulae. Does addition-invariant FO have the same expressive power as FO_{card} on bounded tree-depth structures?

Acknowledgements We want to thank Isolde Adler for bringing the first two authors together with the third author, and Nicole Schweikardt for her helpful suggestions.

References

[1] Michael A. Benedikt and Luc Segoufin. Towards a characterization of order-invariant queries over tame graphs. *Journal of Symbolic Logic*, 74(1):pp. 168–186, 2009.

[2] Adam Bouland, Anuj Dawar, and Eryk Kopczynski. On tractable parameterizations of graph isomorphism. In *Proc. IPEC 2012*, pages 218–230, 2012.

[3] Ashok Chandra and David Harel. Structure and complexity of relational queries. *JCSS*, 25(1):pp. 99–128, 1982.

[4] Bruno Courcelle. The monadic second-order logic of graphs v: on closing the gap between definability and recognizability. *Theoretical Computer Science*, 80:pp. 153–202, 1991.

[5] Bruno Courcelle. The monadic second-order logic of graphs x: linear orderings. *Theoretical Computer Science*, 160(1–2):pp. 87–143, 1996.

[6] Bruno Courcelle and Joost Engelfriet. *Graph Structure and Monadic Second-Order Logic – A Language-Theoretic Approach.* Cambridge University Press, 2012.

[7] Kord Eickmeyer, Michael Elberfeld, and Frederik Harwath. Expressivity and succinctness of order-invariant logics on depth-bounded structures. In *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science (MFCS 2014), Part I*, pages 256–266, 2014.

[8] Michael Elberfeld, Martin Grohe, and Till Tantau. Where first-order and monadic second-order logic coincide. In *Proc. LICS 2012*, pages 265–274. IEEE Computer Society, 2012.

[9] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory.* Springer-Verlag, 2006.

[10] Jakub Gajarský and Petr Hliněný. Faster deciding MSO properties of trees of fixed height, and some consequences. In *Proc. FSTTCS 2012*, pages 112–123, 2012.

[11] Martin Grohe and Nicole Schweikardt. The succinctness of first-order logic on linear orders. *Logical Methods in Computer Science*, 1(1:6):pp. 1–25, 2005.

[12] J.M.Howie. *An introduction to semigroup theory.* Academic Press, 1976.

[13] Leonid Libkin. *Elements of Finite Model Theory.* Springer-Verlag, 2004.

[14] J.A. Makowsky. Algorithmic uses of the Feferman–Vaught theorem. *Annals of Pure and Applied Logic*, 126(1–3):pp. 159–213, 2004.

[15] Marvin L. Minsky. *Computation: Finite and Infinite Machines.* Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1967.

[16] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms.* Springer-Verlag Berlin Heidelberg, 2012.

[17] N. Schweikardt and L. Segoufin. Addition-invariant FO and regularity. In *Proc. 25th IEEE Symposium on Logic in Computer Science (LICS'10)*, pages 285–294. IEEE, 2010.

[18] Nicole Schweikardt. A short tutorial on order-invariant first-order logic. In *Proc. CSR 2013*, pages 112–126, 2013.